

EMERGENCE, ANALYSIS AND OPTIMIZATION OF STRUCTURES

Concepts and Strategies across Disciplines

K. Lucas, P. Roosen (Eds.)

Electronic release of a work-in-progress version, to be published as book when completed.

Version May 3, 2007

(Updates will be provided as soon as they become available.)

Contents

Preface	5
1 Introduction	7
2 Transdisciplinary Fundamentals	11
2.1 Theories, Models, Structures	11
2.2 A Simple Example: Structures in Fluids	14
2.3 Real-world structures and structure-generating processes	19
2.4 Modeling Reality	22
2.5 Model Analysis	35
2.6 Mathematical Methods	46
3 Case studies	61
3.1 Production Engineering	63
3.2 Medicine	77
3.3 Multi-Criteriality Management	85
3.4 Molecular Biology	93
3.5 Energy Engineering	111
3.6 Civil Engineering	123
3.7 Logistics	135
3.8 Innovation Management	137
3.9 Bounded Rationality	149
3.10 Sociology	159
3.11 Acoustics	175
4 Transdisciplinary perspective	187
A Literature	189
B Contributing Scientists	209

Preface

In May 2002 a number of about 20 scientists from various disciplines were invited by the Berlin-Brandenburg Academy of Sciences and Humanities (BBAW) to participate in an interdisciplinary workshop on structures and structure generating processes. The site was the beautiful little castle of Blankensee, south of Berlin. The represented disciplines ranged from mathematics and information theory, various fields of engineering, over biochemistry and biology, to the economic and social sciences. All participants presented talks explaining the nature of structures considered in their fields and the associated procedures of analysis.

It soon became evident that the study of structures is indeed a common concern of all disciplines. The motivation as well as the methods of analysis, however, differ considerably. In engineering, the generation of artifacts, of infrastructures or technological processes are of primary interest. Frequently, the analysis aims at an optimization of the structures and the structure generating processes. Usually a mathematical model is set up, and optimization methods are applied to it. Depending on the problem considered, mathematical or heuristic methods are applied, the latter preferably of the type of biology based evolutionary algorithms. In biochemistry, interest is frequently focussed on the structures of molecules, such as proteins or ribonucleic acids. Again, optimal structures can usually be defined. They are synthesized by elaborate experimental methods, but also by computer based evolutionary approaches, similar in general but different in detail from those used in engineering. Quite different is the study and explanation of existing structures in biology and the social sciences. Here the question of optimality does not present itself in an unambiguous manner. Generally, it must be questioned whether biological and social structures, although developing in an evolutionary way, can be considered as leading towards an a-priori defineable goal of optimality. Also, the method of analysis is usually different. Interactions of individuals forming biological or social structures are strongly governed by large statistical deviations or even non-causal dependencies and require, at least partially, approaches from those of the physical, chemical or the technical sciences. Therefore, mathematical approaches are only helpful to a limited extent. Thought and experimental observation are the basis of analysing and documenting structures and structure generating processes in these fields. The economic sciences play a kind of intermediate role as both types of analysis are found.

In view of the many fundamental analogies of structures and structure generating processes in the various fields, and also of the equally fundamental differences, the workshop resulted in a broad conviction that a systematic interdisciplinary investigation of the topic would be of considerable scientific interest. On this basis an interdisciplinary study group was inaugurated with the beginning of 2003 by the Berlin-Brandenburg Academy of Sciences and Humanities. It was instituted for a period of three years, i.e. over the years 2003 to 2005, with the dedication of studying the emergence, analysis and optimization of structures in the various disciplines. This book documents the results of this study, working out the transdisciplinary fundamentals and analysing their operation in a number of case studies. From the results it can be speculated that a new research field may emerge, condensing the various strategies to a common approach towards a better understanding and treatment of structures and structure generating processes in all fields.

It first became necessary to define the understanding of structures and to find a useful classification for them. Here, the fundamental notions of static and dynamic structures, originally proposed in thermodynamics, turned out to provide a largely common basis for understanding and analysis. Static structures are determined by internal interactions of their constituting components only, develop, as time proceeds towards a predetermined or at least implicitly defineable goal and remain unchanged over time when this goal has been reached. On the other side, dynamic structures, additionally influenced by interactions with the environment, are essentially unpredictable and require steady exchange with the environment for their persistence. Many structures, notably the artifacts in the technical and economic sciences, appear to be static at first sight. However, a closer look reveals that their design and operation depends to a large extent on the environment, which makes them assume typical features of dynamical structures. In particular, when looking at long time developments, even artifacts cannot be considered to be teleological. Rather, these goals change in a basically unpredictable way, sometimes abruptly, and with them the related structures, too. In this general view it appears that the common properties of structures are those found in the dynamic structures, such as dependence on external interactions, abrupt changes and intrinsic unpredictability.

The basis of analysis of structures in all sciences are models. So, the modeling of reality as executed in the various fields had to be investigated. Real structures and structure generating processes in some fields are analysed in terms of optimality, at least partially, i.e. restricted in space and time. However, the assessment of optimality is not straightforward, since various conflicting goals contribute to its definition in most applications. Even if these goals are objective in themselves, their relative weights for a global optimality assessment are in most cases subjective. Sometimes these weights change over time due to subjective esteem. Further, any model aiming at representing reality has to cope with various types of uncertainty. Methods to cope with uncertainty are available, ranging from error propagation and stochastic approaches in the technical sciences to the approach of bounded rationality in sociology. Finally, once a model has been set up that maps the input parameters into a performance function, an algorithm has to be used to work out this relationship, either provided by strict mathematics or by some kind of heuristic.

The complex scenario of qualitative and quantitative modeling, and the goal-oriented and statistical approaches to understanding the meliorisation efforts in the various fields of the sciences and humanities define the scope of this book, with the aim to help readers familiarize themselves with methods hitherto unusual in their own fields.

Klaus Lucas
Peter Roosen

Aachen, Spring 2007

Chapter 1

Introduction

We live in a world of structures. In the form of plants and animals, as well as of landscapes and clouds we are surrounded by the beautiful structures of nature (Fig. 1.1). Human civilization



Figure 1.1: Structures in nature: Plants and animals

has added artifacts of different kind of beauty and geometry, from the egyptian pyramids to the structures of modern urban settlements (Fig. 1.2). Modern industrial development has brought



Figure 1.2: Artificial structures: Complex Buildings

to appearance artificial structures in the form of engines and products of various types, such as airplanes, automobiles, power stations, chemical production sites and many more (Fig. 1.3). We start learning to design molecular structures, such as proteins or nucleic acids, for medical applications by experimental methods. We organize our social and economical life in societies, companies and logistical decisions. We invent medical treatment procedures, mathematical algorithms and theories explaining the phenomena of the world and we create structures of art in music, poetry, sculptures and painting (Fig. 1.4). Structure generation, be it of a material or an immaterial nature, is in a process of permanent development, accompanied by decay and resurrection, and, in particular, characterized by a trend to progress and complexity.

What do we know about the principles of the processes creating structures? Since Darwin we believe that the structures of the living nature emerge from ongoing stochastic mutation and selection, leading to an increasing survival and reproduction of the fittest. The same principles can frequently also be observed to operate in processes leading to our artificial structures, no matter at which area we look, i.e. technical structures, structures in the social sciences as well as structures of art and intellect. This evolutionary procedure is, however, occasionally interrupted by a revolutionary invention or new intellectual concept, apparently without any evolutionary roots,



Figure 1.3: Artificial structures: Airplanes and complex chemical engineering plants

which strikes and alters the world and its further evolutionary development.

Given the fact that structure generating processes are studied in almost any scientific field the question arises to what extent they can be reduced to common foundations. In particular, it appears fruitful to investigate whether progress can be made in analyzing and, where applicable, optimizing them by transferring the knowledge accumulated in the various disciplines across the frontiers.

As can be expected, there are differences between the fields with respect to approach, analysis, and solution strategies, but there is much more in common than scientists focussing on disciplinary work normally anticipate. Therefore in Chapter 2 some transdisciplinary fundamentals are discussed. We start from a first look at the structures and structure building processes in simple systems. The principles are abstracted and applied to the structures of the real world. Then we proceed to the problem of modeling reality. In this context the question of optimality is discussed, and general optimization methodologies are scrutinized.

In Chapter 3 a selection of examples illustrates the issues involved. Problems and solution procedures are presented that are typical for the disciplines contributing to this book, ranging from engineering, over natural sciences up to the social and economic sciences. They are not written as in-depth treatments for the specialist. Rather they are presented with the aim to explain area-specific questions and methods of scientific analysis in a way digestible for non-specialists.

Finally, in Chapter 4, we summarize the results and further challenges of this interdisciplinary approach to understanding structures and illustrate the perspectives of a new research field.




Figure 1.4: Artificial structures: Works of art contain, besides material structures, immaterial as well: Meaning, social context and many others, representing immaterial aspects of society.

Chapter 2

Transdisciplinary Fundamentals

The analysis of structures and structure generating processes in the various disciplines rests on the basis of some transdisciplinary fundamentals, more or less common to all fields of application. These fundamentals will be discussed in the present chapter. First we define our understanding of structures and their properties as treated in this book. After that we discuss the various aspects of modeling real world structures in terms of artificial mental concepts, referred to as models. Here we recognize the necessity of introducing simplifications to the real world structures, in order to make them accessible to analysis. We shall then discuss the fundamentals of model analysis, where we reflect on the question of optimality, of constraints in optimum seeking, of conflicting goals, and of objective and subjective quality. We conclude the chapter by discussing mathematical and heuristic methods of search for the optimum.

Concerning optimization there is a fundamental difference in view for technically oriented structure generation processes and those studied in biology and the social sciences. While in biology and the social sciences there is no teleological development, in the technical as well as the economical sciences we frequently analyze structures with the aim of optimizing them. When a single goal can indeed be specified we can apply standard algorithmic methods that offer themselves for this process. However, a closer look usually reveals the existence of not one but of concurring and potentially conflicting goals. Furtheron, since the structures to be discussed here are created to serve the needs of people, the assessment of quality measures is influenced by subjective apprehension. In elucidating all these realistic facts, we find that in the general view there seems to be no predetermined goal, not even in technology and economy, in particular in a long term perspective. On the other side there are evidently predeterminable goals in a partial sense, e.g. in a restricted short time view, in a restricted area of consideration, or in a restricted system of values. This then seems to close the gap to the biological subsystems, which frequently are partially optimized to a remarkable extent, although the full organism is not.

2.1 Theories, Models, Structures

Three notions will frequently appear throughout the book: theory, model and structure. They are usually utilized in a wide range of contexts. None of the words has a universally accepted definition. Their meaning depends on the disciplinary environment in which they are used and on the educational, sociological, or scientific background of the persons employing them. This book cannot and will not make an attempt to provide precise definitions, but rather will explore the meaning of these notions in different contexts, explain their usage by examples, point at differences, and work out common features with the aim to help improve interdisciplinary communication.

Theories. To scientists, a phrase "...theory" signals a particularly well-tested set of scientific ideas for which there is some range of phenomena where the theory gives correct predictions every time it is applied. A theory represents the well-established understanding of a system of objects, mechanisms, or processes. In the technical and natural sciences, it is often formulated in terms of mathematical formulas. These make the intellectual concepts precise that have led to formulating a theory. A theory can never be proven to be complete and final. New discoveries may result in generalizations, like the special theory of relativity or quantum mechanics extending Newton's laws of mechanics to treat processes with velocities approaching the light velocity and phenomena on an atomistic scale, respectively. What system of ideas is called a theory depends on the history and development of a scientific field. While nobody doubts that there is a mature theory of electromagnetism, some may question that psychoanalysis, has a similar scientific status that warrants its denomination as a theory.

Models. Models are employed to describe certain aspects of objects or phenomena without claiming to represent reality. Typical physical models are the models that are constructed of buildings. An architect would build a model of an opera house in quite a different manner than a sound engineer. While the architect would concentrate on the physical appearance a sound engineer would focus on the acoustically relevant aspects of the interior, such that experiments can be made to simulate the sound propagation.

In the technical and natural sciences, a conceptual or abstract model is often phrased in mathematical terms using variables, equations, inequalities, or other logical and quantitative relations. The aim of such mathematical models is to enable an analysis of the modeled object or process within the underlying mathematical theory. Just as in physical modeling, one and the same object can be investigated with different goals by employing different mathematical models. Each of these mathematical models will ignore certain aspects and focus on others. The reason for using a multitude of models is that we are frequently unable to understand the process or object on the whole. Instead, we try to capture certain aspects by partial models. At present, for instance, nobody is able to set up a mathematical model of an opera house representing all aspects important for the construction of such a building at the same time. One mathematical model may reflect the statics of the building, another air conditioning and heating, a third sound propagation, etc. All these models depend on each other, but to reduce complexity to a presently manageable level, special features are singled out while others are kept fixed.

The relation between model and theory is not clear-cut. The word theory is usually employed for "general aspects". Quantum theory and thermodynamics in physics, algebra and probability theory in mathematics, and equilibrium and game theory in economics are well established theories addressing broad ranges of phenomena. One would, however, not speak of a theory of motor vehicles. Cars are modeled, and there is not only one model. There are mathematical and physical models for the aerodynamics, models for stability and crash simulations, models for the engine, one for the catalytic converter, etc. Many theories are employed to model a vehicle properly.

Theories and models are both judged by their ability to predict consequences. Gravitational theory can be employed to compute the trajectory of a space craft flying to Mars very exactly, and various mathematical models for cars have significantly helped to make vehicles safer, to reduce energy consumption, and produce cars more efficiently.

Structures. Although a precise and universal definition is difficult, it seems that we have an intuitive feeling or maybe even an inherent understanding for what structure is. We encounter structures in many different appearances. They may be material ones, like ordered matter, or immaterial ones, like social interrelations or concepts of thought. Scientific problem solving reveals structures that may be characterized as distillates arising from data analysis, general problem knowledge, experience, experiments, contemplation, and in many other appearances. Which structures are recognized and considered important for the specific niche of the real world dealt with and which structures finally emerge as key objects of study strongly depends, though, on the scientific and cultural background of the investigators and on the viewpoints chosen.

Material objects, i.e. structures, as well as concepts and ideas, are interlinked by more or less complex networks. We refer to this phenomenon of a structure being part of a network generally as its interaction with an environment. To illustrate this view, let us consider a simple object, e.g. a window in a house. A window has a certain internal structure as it consists of translucent elements, a frame, an operating handle, and some inner mechanics. On the other hand the window is embedded in the house and, together with the house, in the neighboring environment that may define varying interactions with it. When these interactions are taken into account other properties of the window become important, like its direction with respect by sunshine, its relative position in the room, its size, its operability etc. . Last but not least, in a sociological or psychological context, the window in the house may be seen as a purely abstract functionality connecting the adjacent sides for visible information while at the same time separating them. This last aspect reflects the window as a purely ideational structure. Each view defines a particular net of interactions, that makes the window analyzable and optimizable in a systematic way. So structuring is a fundamental part of any problem-solving process.

We may even assume that the way our brain operates on its sensory input strongly determines the generation of (abstract) structures. Originating from the ancient impulse of making the material environment manageable by ordering and introducing causality, facilitated by the existence of our cerebrum, we extend this method to purely theoretical constructs as well. Only this ordering process makes them manageable with respect to understanding and, sometimes, melioration. Of course, there is history, tradition, experience, there is process in general: We do not determine structures anew in every act of problem-solving. Embedded in our collective memory, in our language, our rules of conduct, social or technical norms, or our scientific knowledge and last but not least in our artefacts there are given results of former acts of structuring, resulting from problem-solving in other, sometimes similar contexts. We use to start working on the base of these given structurations, we recognize a window when we see 'it'. In our everyday practices working with given structures prevails compared to inventing new structures. Therefore we tend to assume that these given structures are part of 'the nature of objects'. However, this essentialist view reveals itself as insufficient as soon as we come across new problems as it is often the case in sciences and humanities. Solving these new problems requires us to establish new views on the case and its context.

The interactions of a structure with its environment have to be carefully observed in any analysis. Usually, we tend to isolate an object from the surrounding world for analytical and abstracting purposes in the sense of defining a system, i.e. the object under consideration, and the environment. Without this isolation an abundance of influences may have to be taken into account when an object is to be modeled and investigated. Let us return to the above window example that we might like to optimize with respect to its energetic behaviour. In a first try, we

limit our point of view to the interior structuring, considering everything else as a non-interacting environment. Accordingly, we model its inner structure only and subject the available parameters defining it to an optimization procedure with respect to the minimum loss of heat. Then it is quite obvious that the window will turn into something small (because the surrounding walls are better isolated), thick (as more material reduces the overall heat conduction), not very translucent (since heat-reflecting panes tend to be dim), and directed southward (in order to accumulate energy gains by direct sunlight). Clearly the 'optimized' object is not the technical solution we had in mind as human beings when we started out optimizing it. Had we, instead, isolated our perspective to just one of the other aspects we would have come to other, but similarly unsatisfying results. So, any one concerned with the assessment of optimality of a window in a house, such as the carpenter, the psychologist, the police and, finally, the housewife cleaning it, will arrive at different results. We conclude that an 'optimal' window can thus only be achieved if the model includes the relevant interactions with its environment. Recognition of these interactions will by itself lead to a multitude of mostly conflicting and subjective property demands as will be treated later in this book. The impossibility of treating the window independent of its interactions with the environment reflects the common observation that there is a wide variety of window concepts in buildings. They arise from the rather individual interactions with the environment along with a plentitude of conflicting and subjective preference settings. So, an analysis of structures and structure generating processes requires understanding of the objects under consideration and also their interactions among themselves as well as with their environment.

2.2 A Simple Example: Structures in Fluids

In order to arrive at fundamental insight and classifications, a first look on structures and structure generating processes is reasonably focused to highly simplified systems, where the objects and the interactions between them are well understood. An example of such systems are fluids. The objects are the molecules, the internal interactions are those between the molecules, and the external interactions those between the fluid and its environment. Due to the simplicity of these interactions important and formal insights can be attained that will be generalized to much more complex, macroscopic structures in other fields of interest. The science providing a universal theory in such systems is thermodynamics. Thermodynamics tells us that there are two fundamentally different types of structures, referred to as static and dynamic.

2.2.1 Static structures

Let us first look at static structures. A prototype of a static structure generated by nature in a fluid is a snowflake (fig. 2.1). Under certain circumstances, such structures are also referred to as structures in equilibrium. Here, equilibrium is that state of a system attained when it is isolated from its environment and when all relaxation processes have come to an end, irrespective of the process conditions leading to it. So, for static structures in general and equilibrium structures in particular, all interactions with the environment are eliminated and we have a final, i.e. not any more changing state.

The global principles leading to 'static structures' are well understood. In an isolated system, the second law of thermodynamics requires, that the entropy of the system under consideration must increase in any real process. The state of maximum entropy, consistent with the constraints of the isolated system, is the equilibrium state and plays the part of a predetermined goal. In that state the rate of entropy production is zero, because all processes have come to an end. Under suitable constraints (but not always) maximum entropy will require the generation of macroscopic




Figure 2.1: The always varying structures of snowflake crystals. Although every crystal differs from all others there a common structuring principles, like the six-fold geometry.

static structures, such as that of a snowflake. When we follow the structure generating process of a snowflake by looking at the modeled interactions of the molecules on a computer over time, we see that the evolutive relaxation process towards maximum entropy indeed leads to the emergence of a static structure, when the constraints of the system are appropriate. The particular and complex shapes of the snowflakes are determined by the interactions between the molecules, but also on additional parameters such as the rate of cooling. As such, they are not really equilibrium structures but rather structures frozen on the way to equilibrium. As a very simple example for a true equilibrium structure we consider the vapor-liquid equilibrium of a pure fluid. When we cool a pure gas, say water vapor, we observe that the molecules will at a certain temperature generate a structure, in particular a spatial separation of two phases, gas and liquid. This structure, much simpler than that of a snowflake, will persist after isolation of the system, i.e. after eliminating all interactions with the environment, is a static one.

More complicated static structures are known to appear in fluid mixtures as a result of interactions between the molecules by cooling processes from the disordered gaseous state. Well-known examples are the phase and reaction equilibria exploited in an industrial scale in chemical processes. The particular appearance of the structures depends on the particular intermolecular interactions. The plentitude of internal interactions between the objects, is the reason for the enormous plentitude of static structures in fluids, that we observe. The entropy maximum principle explains that the existence of a structure is subject to particular constraints. It will disappear when the constraints are violated, e.g. at lower and higher temperatures and pressures, since the entropy maximum will then require a structureless state. We note that for purely mechanical systems, i.e. systems at zero temperature and entropy, the maximum entropy definition for isolated systems transforms into one of minimum energy. So, from the geometry of a simple molecule, as calculated today in a standard way from quantum mechanics by looking for the minimum energy, up to that of a folded protein, as studied in biochemistry, molecular equilibria can be considered as




Figure 2.2: Folded three-dimensional structure of a protein molecule, representing the state of minimal free energy. The structure is displayed in a stereographic representation. If observed with a squinting view the spatial structure appears after a while of eye adaptation.

static structures (fig. 2.2). In general, the free energy has replaced the energy of purely mechanical systems, when entropy effects are not negligible. We note that conditions can occur in which the minimum of (free) energy is not attained by a system on its way to equilibrium since the required activation energy is not available. We then find structures associated with frozen states, depending e.g. on the speed of change like in the case of snow flakes. However, when the hindrance is eliminated, they invariably move to the state of minimum free energy with the emergence of the associated equilibrium structures.

We summarize that static equilibrium structures in simple fluids originate from the particular interactions between their molecules after eliminating any interactions with the environment. They are generated by relaxation processes towards a predetermined goal, which here is maximum entropy, unless hindered by particular constraints. These relaxation processes can be reproduced on a computer, although the final result can in such simple cases more easily and reliably be calculated in one rational step by a direct mathematical optimization routine. Static equilibrium structures are predictable and reproducible. So, given the full definition of the molecular interactions in a thermodynamic system, the equilibrium state can be formally predicted, at least in principle, and the same static structure will be reached independently from the starting condition as well as of the process of the development. Clearly, static structures in fluids are no more than a simple example for much more complicated static structures placed into the society as technological artifacts, such as engines, buildings etc. We shall see, however, that those artifacts share important fundamental properties with those of the static structures of fluids.

2.2.2 Dynamic structures

Static structures, i.e. structures originating from relaxation processes towards a predetermined goal and remaining constant after eliminating all interactions with the environment, are only a small part of the structures in the world. Evidently there are other types of stable structures which are generated and kept in existence by interactions with the environment through a continuous transfer

of energy and matter in open systems. This input of energy and matter stabilizes a state away from equilibrium, and so, these structures are referred to as nonequilibrium or dynamical structures. They appear to us in the beauty of living nature, as well as in form of processes, social networks and in practically all operational appearances of our man-made artifacts (Fig. 2.3). Again the study



Figure 2.3: A bridge as representative of a man-made dynamic structure. Even if it appears static at first glance, the continuous care in terms of inspection and maintenance (interaction with the environment) should be kept in mind that is needed to retain its functionality over a long period of time.

of fluids, on the basis of interactions between simple molecules, is a convenient basis for getting insight into the generating and sustaining processes of such dynamical structures. Two simple and well-known examples may serve to illustrate the point and will then be generalized.

One is the Bénard-instability, studied in hydrodynamics (Fig. 2.4). A horizontal fluid layer between two plates of different temperatures in a constant gravitational field will show a structureless heat conduction phenomenon at small temperature differences. However, at a sufficiently large temperature difference the structureless state becomes unstable. Convection occurs, and the heat




Figure 2.4: Infrared camera view of the free surface temperature field above the onset of the Marangoni-Bénard instability in a 5 mm layer of silicone oil heated from below. The pattern is hexagonal, apart from the influence of lateral walls. (Courtesy of Physical-Chemistry Department of the Faculty of Applied Sciences — ULB)

flow and along with it the entropy production increases. Although accompanied by a production of molecular chaos, this increase of entropy generates regular hydrodynamic patterns, a stable

dynamical structure. So, although the entropy evidently is increased as a whole there are areas in the system which are obviously characterized by a local reduction of entropy, the structured Bénard cells. A non-zero rate of entropy production, kept up by an input of energy, is the basic source of this structure generating process.

The second example refers to dynamical structures generated by a certain kind of chemical reactions, the Belousov-Zhabotinskii reactions [Win72]. When specified chemicals are transferred to a reaction apparatus under suitable circumstances, a regular change of colours and also a migrating spatially coloured pattern can be produced. A similar moving pattern structured by the same principles of an oscillating biochemical process, although with much lower frequency, has been observed in a special mutant form of mice (fig. 2.5). Again, there is a considerable rate




Figure 2.5: Traveling stripes on the skin of a mutant $Foxn1^{tw}$ mouse, with (a and b) showing side and top views of an adult mouse, respectively. (c and d) Schematic drawing of a typical wave movement in an adult mouse. The numbers by each line represent the time course of a wave. The time interval between the numbers is ≈ 30 days. (e) Traveling waves formed by a Belousov-Zhabotinskii (BZ) reaction. The shape of a mouse is superimposed. The outermost wave is the first wave. Intervals between waves are relatively wider for the first a few waves, then become shorter and constant. Experimental BZ waves are made by the standard method described in ref. [Win72]. To trigger two separate centers of concentric waves, two iron grains had been put on the dish. (Courtesy of Noboru Suzuki, Masashi Hirata and Shigeru Kondo)

of entropy production associated with such a reaction process. And again this rate of entropy production is a source of structure generation.

So, in both of the above nonequilibrium structure generating processes, we find that a non zero entropy production rate is associated with the structure generation and preservation. This entropy production is made possible by specific interactions with the environment. The particular type of structure depends on both, the interactions with the environment as well as the internal interactions of the molecules. When looking at the processes of dynamic structure generating in such comparably simple systems, i.e. fluids, in detail some interesting features become visible that can be generalized to more complex situations. So, we observe that in the course of the still rather

simple-structured nearest-neighbor interactions between the molecules long-range, macroscopic structures are created. In the case of the Bénard effect, i.e. in a hydrodynamic system far from equilibrium, small fluctuations, which are damped close to equilibrium, are augmented as a coherent effect and create a new type of macroscopic structure, stabilized by the exchange of energy with the environment. In the Bénard cells the global temperature difference between the lower and upper interface plate is so great that the *local* environment of the interacting molecules is modified in such a way that in a first step of structure evolution small-scale deviations from the global temperature distribution arise by stochastic effects. They modify the environment and the associated interactions. These cumulated effects, taking place in the immediate surrounding, add up to stabilize and extend the self-organizing and stabilizing long-range structures. Similar effects become apparent in the chemical or biological structures referred to above.

As a common feature of dynamic structures we find an unexpected element of randomness, non-reproducibility and non-predictability. The precise and detailed form of the Bénard cells is not predictable and not reproducible. They may have different sizes and rotation directions. The same is true for all dynamical structures of thermodynamic systems, notably those arising in chemical reactions. The theoretical foundation for this non-predictability is the fact that the appearance of dynamical structures is a stability problem. Increasing the deviation from the equilibrium state leads to a point where this state becomes unstable. The point of instability is referred to as a bifurcation point, at which fluctuations of an intrinsic nondeterministic character lead to an unpredictable new state of the system, which is usually chosen out of various possibilities. Once this state is attained the system will further develop in a deterministic way on increasing its distance from equilibrium until a new bifurcation occurs. This new bifurcation is, of course, predetermined by the outcome of the earlier bifurcation, i.e. there is an element of history in structure generating processes far from equilibrium. They thus obtain an evolutionary character. There is no predetermined goal for the structure being obtained in detail. The particular appearance of the structure is coarsely classified by the particular interactions effective within the system and with its environment, but in detail it is the result of chance.

It should be noted, that increasing the driving forces beyond strengths sufficient for (long-range) structure creation will usually lead to increasingly finer substructures by further bifurcations until at some point deterministic chaos production sets in: The dynamically created structures change so fast that they are no more identifiable as such. Accordingly a certain local predictability can be maintained only for short periods, with decreasing exactness in structuring forecast due to accumulating nascence of new structure nuclei. The ordering principle leading to such states remains the same, however.

We note the fundamental difference of the dynamic vs. static structures in simple fluids. The latter are characterized both by a maximum of entropy in an isolated system and a zero entropy production, while the former are stabilized by a non-zero entropy production, accompanied by an inflow and outflow of mass and energy. Also, static structures are determined by the interactions between the molecules only, while dynamic structures are also influenced by the interactions with the environment. So, contrary to static structures, dynamic structures cannot persist after isolation. They need the interactions with the environment. As soon as the energy inflow and outflow, along with the resulting entropy production, are reduced under a critical value, the dynamic structures will collapse to a structureless or statically structured state.

2.3 Real-world structures and structure-generating processes

Let us now generalize the fundamental properties of structures and structure generating processes, as elucidated by a consideration of simple fluids, to the material and immaterial structures of our

society. Are such material structures as an automobile or a bridge static or dynamic structures? The answer evidently depends on the type of analysis. An artificial structure may well be classified as static, when we consider its production from a constrained viewpoint. Although the internal objects interacting may be quite specific for each object under consideration, they will be identifiable. Let us consider an automobile as an example. Here we have wheels, the engine, carriage and a number of electronic and mechanical devices as interacting objects. The predetermined goal for generating the static structure of an automobile may be minimum cost, minimum fuel consumption or others. Indeed, designing processes may in principle be set up which lead to an automobile of predetermined quality, based on the selection and composition of its interacting components. However, if we broaden our viewpoint and consider the production of an automobile with reference to its performance on the market, we have to consider it as a dynamical structure. It is determined by the interactions with the environment such as acceptance by the customer. Further broadening the view to operation we realize that without the maintenance the automobile will eventually be destroyed and turn into structureless heap of rust and plastic. Without considering the interactions with the environment, no successful automobile will emerge from a design process and can be kept in operation. The external interactions with the customers make it depend on a plentitude of objectives and the plentitude of automobiles on the market reflects such interactions with the environment.

The structure generating processes leading to technical structures can be considered as evolutive relaxation processes towards a predetermined goal changing over time by interactions with the environment. So, they have elements of static as well as dynamic structures. Value assessment for automobiles has changed over time the predetermined goals, such as design, safety, fuel consumption, all being summarized as success on the market. A similar example is a power plant process with a varying number of feedwater preheaters. Subjected to an optimization towards maximum efficiency as the practical aim, such a system will eventually approach an idealized Carnot process by adapting an infinite number of feedwater preheaters. However when minimization of investment and fuel costs is considered simultaneously as predetermined goals, quite different power plant processes will arise as static structures from the optimization process. When external interactions with markets are considered additionally the power plant structure generating process attains features of a dynamical structure. So, when analyzed over a long time period, there will not be a predetermined goal for a power plant, since assessments and values change in the society in an unpredictable way. The actual power plant processes will then adapt themselves to the varying short time goals in an evolutionary process. So, we see, not only the intrinsic interactions of the objects but also external interactions in terms of objective functions and value assessments are responsible for the plentitude of structures that we observe.

While structures of the technical world frequently lend themselves to more or less stringent optimization efforts, those of the social and biological domain must take into account a both dominant and limiting factor with respect to practical manageability: the individuum and superstructures of it as a core concept. In the social sciences a most common structural issue of vast importance is what we call social stratification. That is, the as we know unequal distribution of life chances among the members of a society. Income, education, job opportunities, status attributed on the base of recognized social capital — all this adds up to a view of society in the large as structured in various ways. While social stratification is a type of macro structure effected only indirectly by individual acts, we have likewise to deal with meso-level and micro-level structures where the actors' influence on structure-building processes are easier to see: work flow and hierarchies in organizations, sets of relations in peer groups or the microstructures of interaction most often invisible to actors — at least as long as interactions proceed without crisis and on routine grounds.

Structures, as human beings derive them from observation, are not timeless, but at the same

time they are not necessarily progressive in the sense of being directed to growing objective fidelity or complexity. If we think of the Inuit and their detailed and differentiated knowledge about snow and ice, we can easily see both that the ability to differentiate varies with the tasks to perform and that established fine grained structurations can get lost as soon as they lose their relevance for problem-solving: For driving a car on a road in winter it is pointless to distinguish 30 or 40 different types of snow, as the Inuit do. Instead, we just differentiate by a binary distinction: snow or no snow, frozen or not frozen. In domains where the properties of 'snow' are more important, like in winter sports, at least some additional differentiations (powder snow, crusted snow, ...) are maintained even in our modern civilization. It may be deducted that the thoroughness of structuring is, to a certain extent, economized depending on the subjective practical value it has for the performing human being.

In sociology the most important and at the same time the most problematic relation is the link between structure and acting: While in mechanics it can be said that structures have direct causal impact (like the molecular structure of a certain metal that is directly responsible for the ability of the metal to carry a certain weight), in society things are different: processes in the social domain consist of acting and they necessarily involve actors. The often urged layman view, that structural factors cause certain actions, is problematic since in society people act on the base not of structures, objects, problems or other people. They rather act on their interpretation of structures, objects, problems and so on. Though by and large we might say that social structures have an enormous impact on educational success (as documented in various transnational learning and teaching effectivity studies), this is not necessarily true for every single case. Every detailed study of, say, decision making in educational processes would show how actors continuously (though not always consciously) interpret their context, including aspects of status, economic resources, family traditions in order to come to a conclusion about, say, going for another degree or not.

We conclude that most of the real world material and immaterial structures and structure generating processes that we wish to study in this book belong to the class of dynamical structures. All structure generating processes, technical, sociological, biological, are associated with interactions with the environment, accompanied by entropy production. Based on simple balances of energy and entropy, this requires an input of energy with a low entropy content and the capability of the system to export the entropy produced by the structure generating process. This can easily be verified for the earth as a whole. Here, we profit from the influx of solar energy as the driving force for structure generation. Solar radiation has a low entropy content due to the high temperature of the sun. The necessary entropy export is realized by the low temperature heat radiation from the earth to the surrounding space, which carries all the entropy with it. Looking at technical systems confirms the general conclusions. Low entropy energy must be transferred to the system, such as fuel or electricity, to generate and preserve structures such as a running engine or an illumination during night time. For living organisms the low entropy input is contained in the food. High entropy energy must be exported in the form of waste heat or waste material.

The details of structure generation depend on the specific intrinsic interactions of the system as well as on the interactions with the environment. Generally we have to consider rather complicated interactions in analyzing the structures of our civilized world. They are in many cases not really understood and cannot be fully expressed in terms of mathematical functions. For example, generating the structure of a bridge or a power plant requires the consideration of interactions between material properties, technological components, natural laws, and, in particular, man made values of efficiency and desirability. Sociological structures are determined by interactions between human beings. We shall return to the analysis of interactions in Chapter 3, when we discuss structures and structure generating processes in various disciplines. No matter how complicated the considered structure is: As soon as the inflow of energy and the export of entropy are blocked,

all structures will collapse. A bridge without maintenance, i.e. after isolation from all artificial energy and mass transfer, will eventually rust away in a structure destroying process. A structure in logistics will cease to exist when nobody will survey it and keep up its existence. An infrastructure will eventually disappear when it is not continuously used. A structure of thought will be destroyed as soon as the human beings adhering to it will stop thinking it. The analogous fate is associated with any natural and artificial structure, including human life.

It is clear that all structure generating processes in principle share the basic properties of dynamic structures, such as non-reproducibility, non-predictibility in detail and lack of a predetermined goal. This is immediately evident, e.g. in sociological structures, in those of diseases and in those associated with weather and climate. The associated structure generating processes are difficult to control and impossible to predict in detail. However, the laws of their generation can be analyzed by considering the relevant intrinsic interactions and fruitful generalizations, as well as average predictions can be made by statistical methods. Furthermore we will see that even dynamic structures *can* be optimized by deliberately blinding out their dynamical nature in representing them by simplified, quasi-static models. This is a reasonable strategy as long as the human optimizer is aware of its limitations. Accordingly, the intrinsic interactions in such system models as well as those with the environment are to be carefully modeled, with the dynamically dominated effects of bifurcations and deterministic chaos etc. being as precisely accounted for as possible.

2.4 Modeling Reality

It is an established fact that the vast majority of systems or processes in the real world are so complicated that there is no hope and even no sense in trying to analyze them in full detail. Instead, scientific analysis has to be liberated from the confusing plentitude of phenomena in order to reduce the complexity of seemingly unsurveyable problems to something that is amenable to analysis. The method of analysis may be mathematics in the technology and natural science oriented applications or observation and thought along with creating notions and their operational interactions in the social sciences and in the humanities. The very process of modeling even a small part of reality is naturally accompanied by a loss of realism, in the sense that some aspects are deliberately eliminated from further consideration. So, there is always the danger of an unacceptable disparity between model and reality, which has to be taken into account when conclusions about model behavior are transferred to conclusions about reality. Models cannot be justified or evaluated within the categories right or wrong. Rather, they are either useful or not useful. The judgement about this classification is subject to a comparison with reality, e.g. by experiment in the natural sciences.

2.4.1 State Models

We wish to illustrate the problem associated with modeling reality first with relation to state models. Under 'state models' we summarize models representing a well-defined state of a system. We start with a simple example, once again chosen from the thermodynamics of fluids. We shall treat varying levels of complexity, each level being adequate to describe certain physical phenomena in different state domains of a substance. We know, that almost any material can be encountered in a liquid or a gaseous state, both being summarized as the fluid state. How can the



Figure 2.6: Fluid phase behaviour of Ar, shown by its molar isobaric heat capacity divided by the universal gas constant against temperature, for a limited range of pressure and temperature.[LUC1]

properties of this technically very important state adequately be modeled, e.g. in order to derive information about its behaviour in certain artificial boundary conditions, like chemical equipment?

Let us first consider a gas of monatomic molecules at normal temperatures and pressures. An adequate model for this reality is that of small billiard balls. This implies a particular conception of the nature of their interactions. In the billiard ball model of a fluid the members are hard balls with a negligible volume, and the nature of the interactions is such that, apart from occasional elastic bounces, there is no interaction at all. So, the billiard balls move independently of each other and their motion is described by the laws of classical mechanics. On the basis of this model one finds formal expressions for the energy, the temperature, the pressure and the entropy of such a gas which are in perfect agreement with the properties of a real gas like Argon over a significant temperature range at normal pressures. To illustrate, we plot the molar isobaric heat capacity of Argon divided by the universal gas constant against temperature in Fig. 2.6. It is a constant of value $5/2$, in perfect agreement with precise measurements. This quantity determines the temperature dependence of thermodynamic behavior and plays a significant role in various engineering applications.

It is remarkable that such an ultrasimplified model can be useful in describing the properties of a fluid. It claims that the molecules of Argon do not have any type of interaction with each other, neither that of attraction nor that of repulsion. Quite evidently this is wrong since it is in conflict with our basic knowledge of the interactions between the molecules in fluids. The contradiction is only apparent, however, and its resolution is that the simple billiard ball model is only probed in a very limited region of the phase diagram, i.e. only at normal temperatures and pressures. In this region, indeed, the interactions between the molecules, while basically effective, are not essential for the macroscopic properties of the fluid, due to the relatively large average distance between them. Events, in which two molecules approach each other so closely that the interactions between them become noticeable are simply too rare to be statistically significant and therefore can be neglected in calculations of the macroscopic properties. So, while the billiard ball model is a useful model for the properties of a real monatomic gas at normal temperatures and pressures, it does definitely not represent the general situation in an arbitrary fluid.

When the properties of a fluid like Argon are to be analyzed over a larger region of states, it becomes necessary to change the model. Going to higher temperatures, say 10,000 K, real Argon starts to ionize, an effect which cannot properly be described by the billiard ball model since it does not allow for a disintegration of the balls into nuclei and electrons as happens in reality.



Figure 2.7: Fluid phase behaviour of Ar over a large region of states.[LUC2]

Going to higher pressures and lower temperatures the phase behavior of Argon becomes entirely different from the prediction of the billiard ball model. Fig. 2.7 shows the fluid phase behavior of Argon over a large region of states. The region adequately described by the billiard ball model is shown as a shaded region in the phase diagram.

A most significant deviation from the prediction of the billiard ball model is the effect of condensation. This is a simple example for a structure generating process in a fluid. Increasing the pressure at a sufficiently low temperature from the low pressure billiard ball region eventually leads to a situation where the average distance between the molecules decreases so strongly that the short-range interactions cannot be overlooked any more. Condensation takes place, and in a gravitational field the liquid and the gaseous region are separated by a meniscus, clearly pointing to the appearance of a spatial structure. Hence, a model for a fluid over the full region of states requires the interactions between the molecules to be taken into account. These can be modeled on the basis of quantum mechanics. With these interactions the properties of fluid Argon can then be predicted over a large region of states by the methods of statistical mechanics along with mathematical algorithms. We note that even this more demanding problem of predicting the fluid phase behavior of Argon over a large region of states, including the structure of the vapor-liquid equilibrium, is usually based on a model which significantly reduces the real complexity. This particularly holds for more complicated molecules, where the interactions depend on their internal structures, and so, on many coordinates. The models for the intermolecular interactions currently used for predicting structures of fluid phase behavior are strong simplifications of the real world. However, when used in combination with an adequate model for the connection between the macroscopic fluid phase behavior and the intermolecular interactions we are able to make predictions of great practical value. This is shown in Fig. 2.8 where a prediction of the vapor-liquid equilibrium in the system is compared to experimental data. Evidently, the model, while being far away from the true situation in the fluid, is a useful representation of reality in the context of vapor-liquid equilibrium prediction. It is by no means guaranteed, however, that this model predicts equally well other aspects of fluid phase behavior of the same system, as can indeed be verified.



Figure 2.8: Vapor-liquid equilibrium modeling prediction compared to measured data for Benzene-Hexane, showing good conformance.[LUC1]

Various aspects of static models, which hold quite generally in all branches of science, can be studied from the simple example of fluid phase behavior. First, a model has to be tailored to the problem to be discussed, e.g. predicting the structure of the vapor-liquid equilibrium in a fluid mixture. There is no sense in extending the model to more and more complexity if only a segment of reality is supposed to be analyzed. So, when the gas phase properties of a fluid such as Argon are to be analyzed, there is no sense in setting up a model for the intermolecular interactions between its molecules. Further, when only the vapor-liquid equilibrium structure in a fluid mixture is to be considered, a satisfying model can be based on a rather crude model for the intermolecular interactions. The demands on model complexity become more and more complex as further aspects of fluid phase behavior are to be studied, such as liquid-liquid equilibrium structures and heat effects. A model representing a large segment of reality will necessarily be much more complex and thus require more effort on applying it to the prediction of fluid phase behavior by statistical mechanical and mathematical methods.

Models for fluids are a simple example for static modeling. Many aspects of them can be generalized to more complicated static structures. Generally we shall, with increasing demands of reflecting reality, need an increasing detailedness of the model. As an illustration, let us have a closer look on the energetic supply of a residential area with district heating and electrical power supply. We treat the system as isolated from the environment. The energy supply system thus represents a static structure. A first order of rather crude detailedness is the consideration of the added-up peak needs in heating and electrical power for the intended set of houses. If the existing electrical power generators and heating stations provide these respective maximum values the demand will be satisfied. If the district is to expand by building additional houses and a sensible statement is required on how many of them could be heated with the same system without increasing power, the power requirement model has to be refined. The coincidence that all customers are drawing their peak requirements at the same time will never occur. Hence it would be much more realistic to model the requirements by a time-resolved demand structure on

known peak demand days separately for electrical energy and heat. An efficient energy system for the residential area could be based on cogeneration technology which requires a correlation of electrical power and heat demands, opening up the possibility to utilize low-temperature waste heat of the electrical power generation for domestic heating purposes. Here the coincidence of respective demands of the two types of energy will have to be mapped, requiring again a more complicated model.

2.4.2 Process Models

In the context of analyzing structure generating processes further aspects of modeling arise. In the technical and natural sciences we here explicitly base our considerations on process models with a certain self-organizing intelligence. Such a model is particularly fruitful in our interdisciplinary context since in many applications it can be implemented in the form of a computer code with an ab-initio unfathomable abundance of possible system reactions upon changes of boundary conditions. Such models are usually called simulation systems, as they try to mimic the causalities of the real world. They are models for structure generating processes.

In the humanities a somewhat different notion of process modeling prevails, as the primary goal of those disciplines is mostly description, not targeted change. Here a descriptive social network is one of the primary ways of modeling, contrary to the simulation of physical causalities. Through the use of formal representations — like directed graphs showing how each unit is linked to certain other units in a network — the model-based analysis of social networks has the aim of uncovering structural patterns and studying their emergence, their consequences, and their temporal transformations. Various techniques are employed to develop respective models as well as to analyze them using mathematical, statistical, and computational methods (e.g. simulation of network processes). In sociology, archives of network data and related computer programs to analyze them are typical resources available to social network analysts. In this way, the area of social network modeling relates to the general goal of understanding how social processes generate, maintain and change social structures, and how social structures both serve to constrain and enable social action. Typical modeling target issues are structured social inequalities, including social demography, socio-historical research, discourse analysis, ethnographies, and policy analysis.

In the causality-emphasizing models a given set of starting conditions is fed into the implemented simulating rule set of mutual dependencies, and the computer calculates the evolutive development of related output or monitoring variables. Even if only a limited number of causal dependencies is implemented in the computer code the variation span of possible outcomes can be vast, due to the effect of the so-called combinatorial explosion. This is due to the fact, that possible outcomes of modeling variable state changes tend to be multiplicative. A nice episode in the musical domain, taken from [Ste99], pinpoints this phenomenon:

„As a boy, John Stuart Mill was alarmed to deduce that the finite number of musical notes, together with the maximum practical length of a musical piece, meant that the world would soon run out of melodies. At the time he sank into this melancholy, Brahms, Tchaikovsky, and Rachmaninoff had not yet been born, to say nothing of the entire genres of ragtime, jazz, Broadway musicals, blues, country and western, rock and roll, samba, reggae, and punk. We are unlikely to have a melody shortage anytime soon because music is a combinatorial system. If each note of a melody can be selected from, say, eight notes on average, there are 64 pairs of notes, 512 motifs of three notes, 4,096 phrases of four notes, and so on, multiplying out to trillions and trillions of musical pieces.“

So, even if the amount of modeled variations in an evolutive simulation system are few and easily manageable, the potentially reachable variation space need not necessarily be so: *Simple rules can lead to complex and seemingly purposeful behavior.*

Proof of this fact can be given and supported by more or less abstract formalisms, but an illustration may be more convincing. For this purpose we choose a two-dimensional cellular automaton, known as John Conway's game of life, running on a rectangular grid of infinite or finite size. Grid nodes (mostly depicted as cells) are occupied or empty, and there are only four rules determining the pattern in the next generation:

- i) Occupied cells with no or one occupied cell in the neighborhood become empty,
- ii) occupied cells with two or three occupied cells in the neighborhood stay occupied,
- iii) occupied cells with four or more occupied cells in the neighborhood are emptied, and
- iv) empty cells with three occupied neighboring cells become occupied.

In this concept the 'neighborhood' is defined as the set of eight cells around the one in question.

Given these rules the unfolding of the cellular automation in time is completely determined by its initial conditions, but nevertheless, an incredible richness of different dynamical behaviors results from the 'game of life' rules. Even very slight changes in the pattern of initially populated cells can change the development of later generations dramatically. Some of these patterns are even suggestive of purposeful design. For example, from one neither random nor fully ordered initial pattern a forever living periodical structure develops that emits small gliding motifs in one precisely defined direction. Figure 2.9 shows the recurring sequence of populated cells' distribution that is reached after some time of more or less irregular approach to it. As anticipated, larger structures, although adhering to the same simple rules, show even much more diverse and astonishing features up to the point of reacting to 'environmental' influences in almost intelligent manners [Hen06].

Contrary to the Game of Life, characterized by a very small rule set, a well-known evolutive simulation engine, the computer game 'Civilization' [Wika] (fig. 2.10), may serve as an example that combines a large rule set and a mixture of continuous and discrete variables. It is a model of society development. In a vast network of mutual dependencies human activities like exploration of unknown territory or technology, war and diplomacy are simulated in their effect on an endeavoring society. The player, thought as a ruler of his people, has to make decisions about which improvements or units to build in each city, where to build new cities, and how to transform the land surrounding the cities for maximum benefit. The number of possible simulation engine reactions is so large that it can be interpreted as almost limitless, even though there are boundary conditions for every variable. The actual state of the game is represented by a larger number of system variable states. Apart from the variety of purely deterministic simulation answers there is an additional inclusion of stochastic elements, like a modeling of natural catastrophies. These set of features causes each game to develop different from any preceding, even if the player tries to repeat a former scenario.

Similar simulation systems, although less popular, are used in scientific modeling of development processes, e.g. to understand the development of social groups depending on political boundary condition settings over time. As soon as the behaviour of human beings is modeled the simulations tend to become unrealistic, though: Representing the twists of the human mind in its reaction to more or less comprehensible boundary conditions is still a wide field of future simulation development.

A final remark appears adequate in relation to such process simulation in systems. Mathematical analysis strongly favors steady and differentiable dependencies, since they allow the application of rigorous and well-assessed optimization algorithms for system melioration. But this represents an undue simplification. Instead, typical simulation systems also contain large numbers of

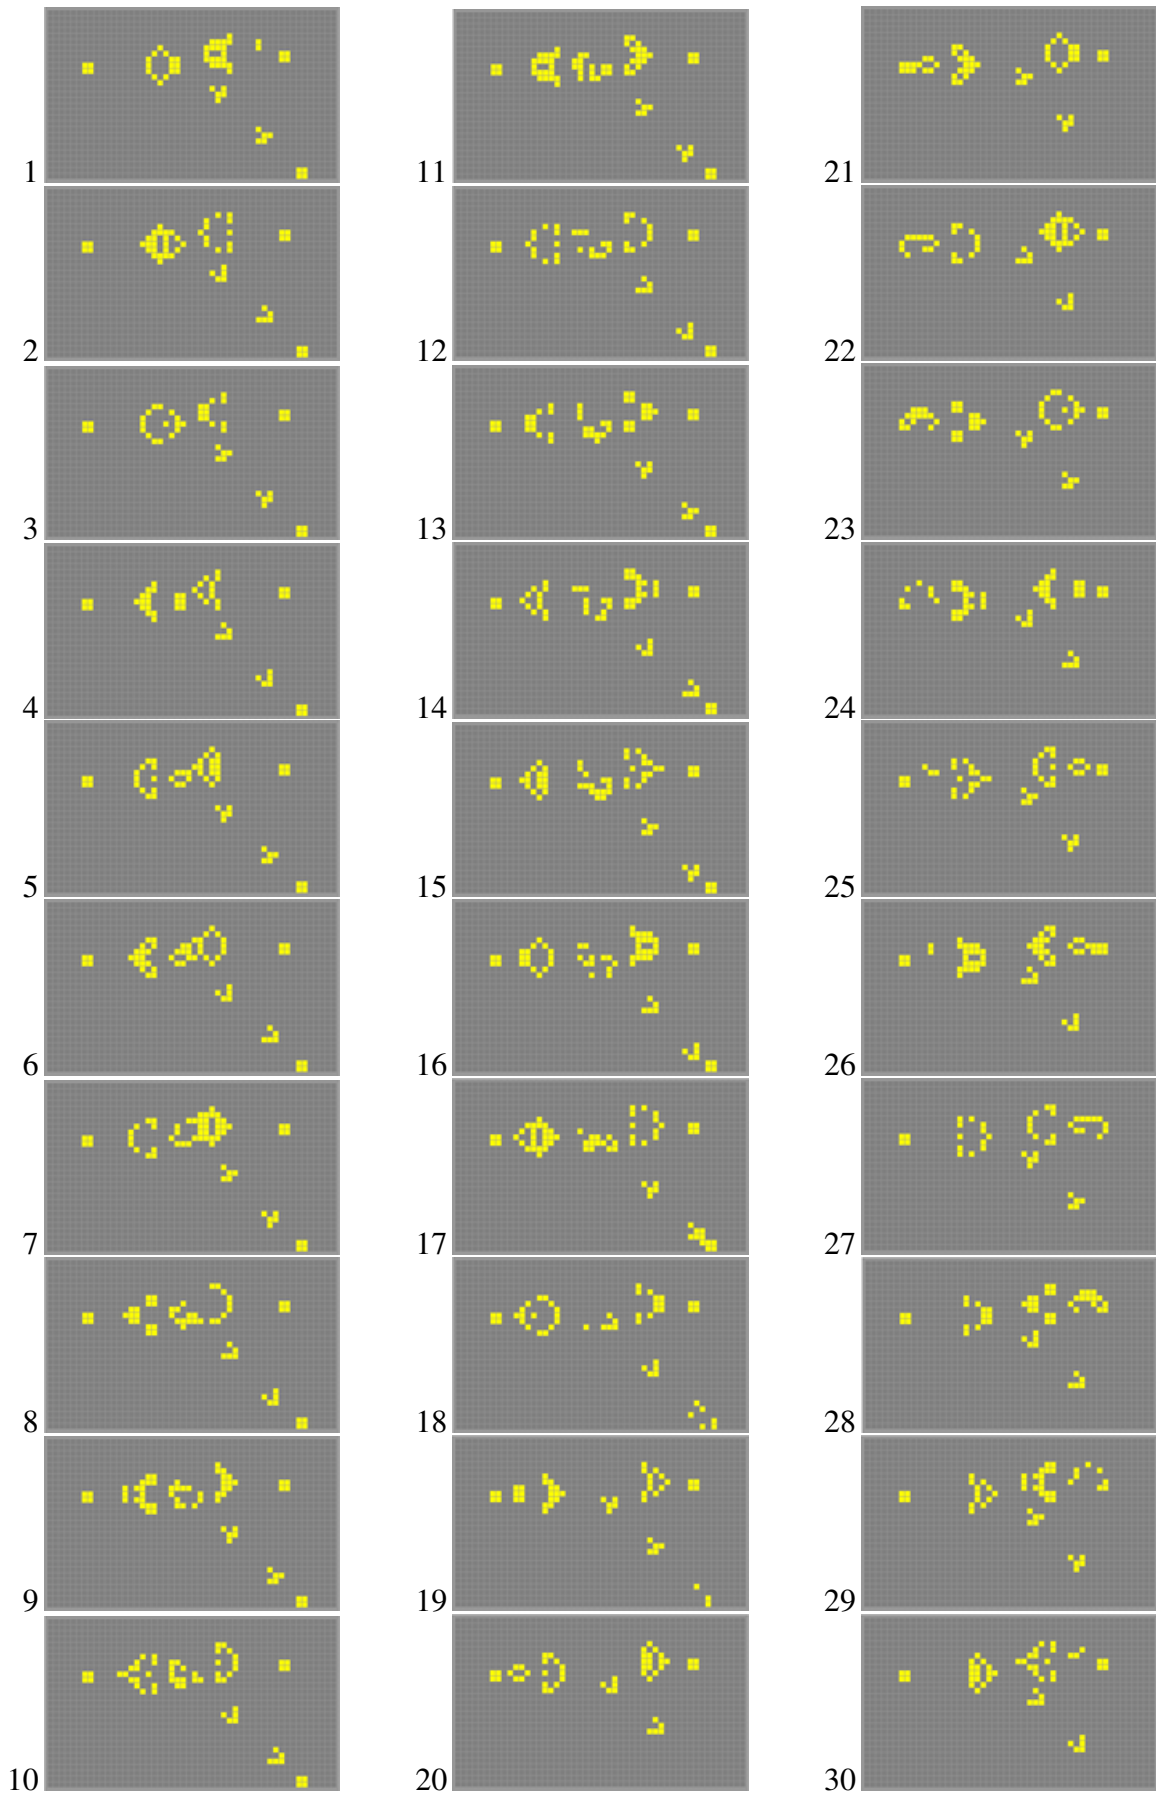


Figure 2.9: Example sequence of a Game of Life configuration, repeating its core structure with a period of 30 iterations and continuously producing 'gliders', being emitted to the lower right corner. Images are sorted column-wise.



Figure 2.10: Screen shot of the society simulation game 'Civilization', taken from [Wika]. Some of the state-determining variables (like amount of food for the people, top left, or treasury of state, bottom right) are depicted or listed in the corners of the screen. Others, like the number of road and hence the facilities of transportation, are represented by respective parts of the drawing.

internal decision variables of a discrete nature to reflect real-world conditions: If, for example, the real features of a hot water piping in a house, or the set of interconnections of chemical engineering apparatuses, is to be realistically represented, the discrete diameters of tubes available for construction purposes must be respected. If any calculation in the respective model is dependent on tube diameters its outcome will accordingly 'jump'. To represent tube diameters in such systems with a continuous variation usually impedes realistic calculations — especially so, if optimality points are sought in a respective framework.

2.4.3 The Economy of Modeling

Connected to the problem of an appropriate model complexity, discussed from the perspective of real-world reproduction fidelity, is the consideration on the effort that should go into the conception of a model from a practical point of view. Even if more detailedness would be desirable in many cases of analyzing efforts with respect to a noticeable reproduction fidelity improvement, it may prove rather impractical to do so. Two main reasons to restrict modeling depth exist: i) the time requirements for evaluating a complex model, and ii) the required labour to set up the complex model, compared to the expected improvement the model promises to cause in an analyzing project.

An example for the first case is the weather forecast. If there is a model to describe a complex phenomenon like the atmospheric development, but the time needed to converge the modeling equations is too large, the real phenomena that should have been forecasted are on a smaller timescale than the forecast itself. The results of the model is then useless. This effect of not satisfied computing power requirements is in most cases a transient one, caused by transient non-availabilities of adequate calculational power. Similar effects show up if rather fast phenomena, like real-time control units and similar applications, are to be scrutinized. The operation of an

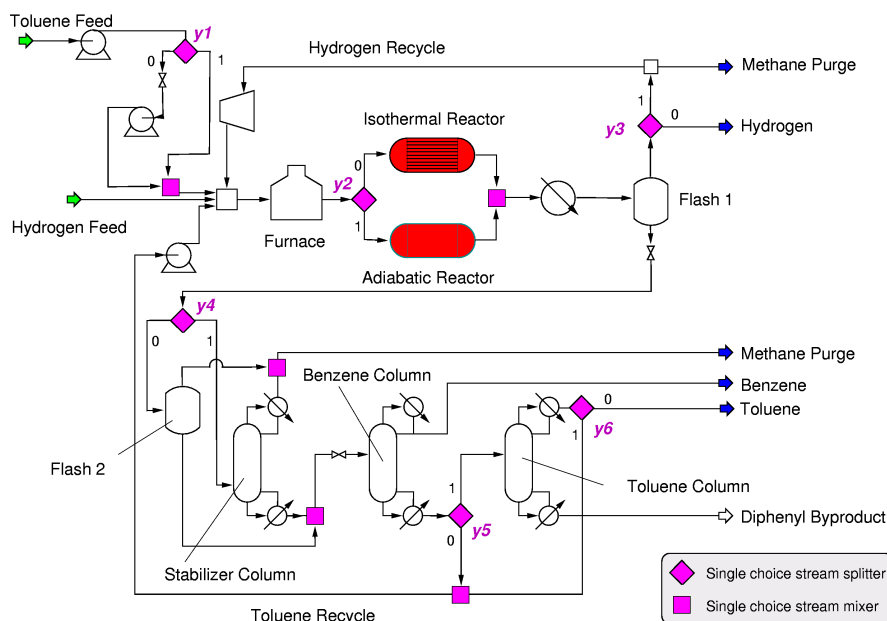


Figure 2.11: Flowsheet of a benzene production plant, taken from [GR99]. Two educts are fed into a (mostly) serial connection of several basic apparatuses (reactors, pumps, distillation columns, . . .), producing the desired output of benzene, but possibly an additional by-product (diphenyl), depending on the concrete choice of the chosen path.

unmanned vehicle, striving to cover a given distance in a sensible amount of time, is such an application. Equipped with a set of sensors, like stereographic cameras, tire slip and bearing controls etc. the steering control unit must interpret many signals in combination, as no single input channel provides satisfying information on its own. Each channel input is ambiguous to a certain extent, so the interpretation of all available inputs needs to be balanced. If the model underlying this interpretation is too complex and thus too slow the vehicle is stuck in the ditch before the control unit gives the signal to circumvent the obstacle. Here a timely circumvention of a non-existing ditch, assumed by a crude but conservative interpretation of the acquired sensors data, is better than a too late one of a correctly recognized ditch.

The economization of the required human power input into the modeling of a given custom case may serve as an illustration for the second case, being discussed on the basis of a typical problem of chemical engineering problem. A general task in this field is the anticipatory performance evaluation of a coupling of several basic apparatuses into a so-called flowsheet (fig. 2.11).

The boundary conditions the real plant has to obey are manifold, being mostly of thermodynamic nature. First of all, conservation laws for matter and energy have to be fulfilled. Second, the properties of the reacting substances, like separation efficiencies in distillation columns, impose limits on what a single apparatus can achieve. Third, at several points in the flowsheet there are decision points on the principal layout of the future plant. In the flow sheet this is represented by alternative paths the materials can take, leading to different product strategies, (human) reaction possibilities upon changes in market prices, apparatus developments etc. . Fourth, there are up to two recycle branches, recirculating some intermediately created substances back to the plant inlet. These potential recycle streams must be considered in the balances.

Basically there are two differing modeling approaches to obtain an optimal path choice and setting of respective operative values — the equation-based and the sequential modular approach. In an equation based approach all simultaneously valid equations and inequality conditions are formulated for the complete system, creating quite a vast mathematically described collection of

dependencies. The problem class may generally be defined as a mixed-integer non-linear (MINLP). Note that in this approach values and settings expressing conditions in the educt section of the plant are frequently connected to values at the end or in the center of the system. Accordingly it is a very complex and time-consuming task to model a chemical reaction system on this basis. Generally adequate algorithms are not available. If a particular solution procedure has been tailored to a given problem, one is rewarded by a 'closed model' that can be subjected to a purely mathematical, rigorous and usually very effective and fast treatment. If small parts of the setup are changed, a lot of reorganizing work has to be put into the model, though, due to the strong interconnections even of spatially far separated units.

There is a caveat in the depicted flowsheet, though. It contains six decision points where a binary selection of one path or the other has to be set. Depending on the chosen path, the mathematical system of equations and inequalities changes significantly, barring a combined modeling of the complete plant. Instead, separate models must be created for every sensible combination of binary decisions. Even if some of the $2^5 = 32$ purely combinatorically defined ones can be dismissed by engineering reasoning quite a number will have to be elaborated to cover the interesting range. Accordingly such a modeling treatment is very expensive with respect to the involved human effort.

An alternative method is the setup of a sequentially modular model where mathematical representations of the individual apparatuses are coupled in a manner strongly resembling their physical interconnection: Substance and energy streams are coupled mostly on a local basis, splitting up the very complex interrelationships (like: the sum of discrete mass flows summarized over every parallel stream of the flow sheet must remain constant) into local subsystems called unit operations. These unit operations, being internally represented by smaller sets of equations and inequalities, may relatively simply be switched on and off for decision variable resetting, requiring rather few changes in global balance adaptations. The modeling effort on the whole is less demanding than the equation-based one, but produces, by its very nature, a computational problem. Due to the localized, hierarchical nature of the modeled interactions the unit operations need to converge in themselves before respective stream redefinitions are to be forwarded to the hierarchically higher level of the unit-combining equation sets. Especially in the case of recycling streams such models have a tendency to oscillate as changed output streams of one unit influence the next in row, but eventually affect their own inputs again due to the circular loops. Such computational problems *may* be overcome by some increased calculational efforts, like decreased time steps, or upper bounds on the changes of important variables with auto-adapted time step decrement.

To summarize, the setup of an equation-based simulation mode requires an elevated amount of human effort that will be honored both by fast optimization times and, even more important, more stable system answers in case of mathematically difficult conditions. Whether this effort is useful in a practical sense depends on the expected return relative to an existing design. Similar decisions on the sensibility of respective modeling efforts are commonplace in most meliorating problems, up to the point that an elaborated treatment may not be adequate.

2.4.4 Modeling Uncertainty

In many cases the real world shows numerous imponderabilities which pinpoint different categories of uncertainty. Appropriately embodying these uncertainties into engineering models provides ways and means to express the limited rationality of a plentitude of real world phenomena computationally. In harmony to the definition of uncertainty introduced by BOTHE [BOTHE93],

uncertainty can be understood as a gradual assessment of the truth content of a proposition, related to the appearance of a specified event.

The following diagram (Fig. 2.12) demonstrates the potential subcategories of uncertainty and assigns them to corresponding theories that have been elaborated in the last two decades. Three different subcategories can be distinguished:

- stochastic uncertainty, which describes the random results of multiply repeated experiments where the boundary conditions have always to remain unaltered,
- informal uncertainty, which describes information deficits due to the limited information sources and an only small number of observations,
- lexical information, which quantifies a relevant real world fact in terms of a linguistic variable by ranking the membership of that fact to a defined uncertain set.

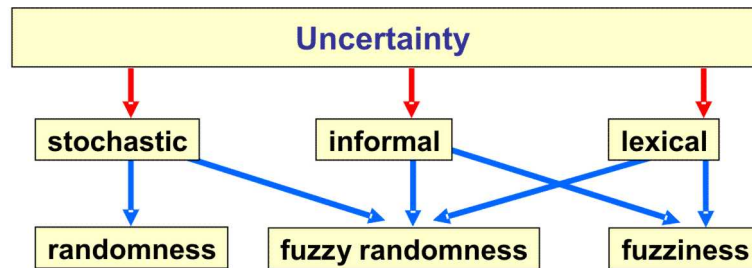


Figure 2.12: Subcategories of uncertainty

There is a rather traditional treatment framework to calculate the effect of uncertain input or measurement values on the outcome of a functionally dependent output value: the error propagation method. This kind of uncertainty effect is almost omnipresent in the modeling of real world systems. It poses no major problem on interpretation of results as long as the model is not a self-evolutive one, meaning that the modeled responses of the system are directly depending on a set of input parameters. As an illustration, let us consider a very simple model for determining the required power output of a vehicle engine, driving at constant speed on a motor way. The modeling is done with a simple formula:

$$P = \frac{\rho}{2}(v + v_0)^2 Av c_w$$

(with ρ = density of the air, v = velocity of the vehicle, v_0 = headwind velocity, c_w air resistance coefficient, and A = head face of the vehicle)

The rule of error propagation then directly tells us how uncertainties in a given variable, say, the air resistance coefficient of the vehicle c_w , lead to a respective uncertainty in the calculated required locomotion power:

$$\Delta P = \frac{\rho}{2}(v + v_0)^2 Av \cdot \Delta c_w$$

Or, with other words, the *relative* uncertainty of required power, $\Delta P/P$, is identical to the relative uncertainty of the vehicle's air resistance coefficient, and the *absolute* uncertainty in P is constant for an assumed constant uncertainty in c_w . This is a dependency the common mind usually expects. On the contrary, if we apply calculus to the modeling equation to elucidate the effect of uncertainty in the velocity value, we find for the absolute error in P

$$\Delta P = \frac{\rho}{2} c_w A (3v^2 + 4vv_0 + v_0^2) \cdot \Delta v$$

or, if we again look at relative uncertainties,

$$\frac{\Delta P}{P} = \frac{3v + v_0}{v + v_0} \cdot \frac{\Delta v}{v}$$

which gives us a definitely more complicated dependence to reflect on. If we experience no headwind ($v_0 = 0$) the relative uncertainty in P is three times as large as that in our observed variable v . This dependency can change dramatically, though, if we have a stronger following wind ($v_0 < 0$), e.g. in the order of half the speed over ground. In that case, a miss on our measured velocity will raise the propagation of relative uncertainty to a factor of five, far away from what the common mind would expect!

Although this effect of error propagation may come unexpected for a person not regularly involved in technical calculations, things tend to grow much worse in case of self-evolutive models. In these models, frequently applied to simulate the temporal development of a system, the calculative result of a value at a certain time interval Δt_i is needed to determine its value at then next interval Δt_{i+1} and later stages. In such cases simple error propagation techniques will soon lead to accumulated errors well beyond the actual parameter mean values, thus rendering their calculation worthless. In such cases, more elaborated techniques, based on in-depth probability and density function considerations, must be put to work.

Well known since years, but still advancing, is the probabilistic approach or the concept of randomness which captures uncertainty through random variables and/or random processes in an objective fashion, excluding subjective views on the problem. Thus, based on long-lasting observations or experiments, an effective probabilistic assessment of the quantities, patterns, actions and structural responses within engineering systems and system processes can be made available. For stochastic variables, the key ingredients are expected values, mean values, variances, and also moments of higher order, quantile values, probability density functions (pdf) and cumulative density functions (cdf). For the wide variety of time-variant stochastic processes, functions for the average behavior, auto-correlations, auto-covariances, covariances and auto-correlations as well as cross-correlations are of interest. All of them can appropriately map the stochastic phenomena of uncertainty in engineering scenarios. Two problems, however, have to be realized if conventional randomness is applied in reliability analysis or optimal reliability-based optimum design: (i) a sufficiently large sample size is required if the pdf and cdf of a random variable have to be accurately defined, (ii) often, it is not an easy task to find the characteristic description of a real world process, with respect e.g. to stationarity, ergodicity, spectral moments, etc.. Consequently, if unfounded assumptions are made when a stochastic simulation model of a technical system or process is created, then additional uncertainties are induced on top of the attempted incorporations of uncertainty into a model (i.e. uncertainty of the uncertainty).

In the past years, fuzziness has tremendously increased in prosperity as a new paradigm to differently computerize uncertainty in engineering. Introduced already by ZADEH [ZADE65] around 1965 despite numerous hostilities, the fuzzy-based optimal decision finding and fuzzy control have been extremely successful in many engineering applications. In particular, this applies for highly complex problems of automatic control engineering where the solution of non-linear partial differential equations represented the regular *modus operandi* for a long time. As a matter of fact, many highly non-linear or near-chaotic problems, which could not appropriately be solved through numerical simulation (e.g. by means of non-linear finite element models), proved manageable by applying the fuzziness paradigm. The nucleus for fuzziness is the human attitude, even if no viable numerical representation of a problem is possible, to express solution behaviors

in terms of fuzzy rules which connect uncertain input with uncertain output quantities. Both, input and output quantities are modeling informal and lexical uncertainties, i.e. non-statistical properties, where particularly subjective aspects are entered into the consideration. To this end, specified crisp input and output quantities of a basic set are transformed by fuzzification into fuzzy sets using so-called membership functions of different functional shapes (e.g. triangular, trapezoidal or curved). Based on the above tripod 'fuzzification', 'fuzzy rules and defuzzification', 'fuzzy in fuzzy control', and based on the α -discretization and α -level optimization in structural analysis (see book published by Moeller [MOELL04]), data as well as model uncertainty of engineering systems and processes can be verifiably be handled. It should be mentioned that fuzziness can analogously be expanded from elementary quantities to functions leading to fuzzy functions, fuzzy processes and fuzzy fields.

Fuzzy randomness has to be introduced if neither randomness nor fuzziness is sufficient to describe the 'crude reality'. This happens when the rigorous preconditions and laws of randomness can not be matched, e.g. because a stochastic quantity is affected by informal and/or lexical uncertainties. Typical for this are (i) only small numbers of samples such that the type of the pdf or cdf has to be approximated with considerable uncertainty and (ii) the violation of fundamental principal of probability, i.e. the constant reproduction condition, by which it must be guaranteed that all samples of the taken universe obey identical boundary conditions. In practice, both premises are often violated. The universe may contain only a few samples, in many cases the boundary conditions are varying with respect to time and location. Also, the assumption of statistically independent stochastic variables and/or stochastic processes, used in many engineering application, does not correspond with real world scenarios. The effect of how a random quantity is fuzzificated can be seen in (Fig. 2.13) where a cdf is superponed with fuzziness.




Figure 2.13: Fuzzification of a Stochastic

In view of these unavoidable uncertainties all model predictions necessarily assume the character of a range of values, depending on the sensitivity of the parameter uncertainties to the macroscopic predictions. Any rational application of a model, including forecasts on the optimal target function values and settings of parameters, must therefore take such uncertainties and their consequences on the conclusions into account. This observation affects in particular one typical method of model enhancement: the inclusion of increasing numbers of system describing parameters. On the first glance, the addition of more details will improve the detailedness of the model, providing more adjusting screws to fit it to the observed reality. But this holds only as long as the additional

parameters can be input with sufficient sharpness. If models tend to be unsharp, however, additional parameters will only result in an obfuscation of the simulated causal dependencies.

2.5 Model Analysis

We now proceed to the fundamentals of quantitative analysis of a model. The motivation for this analysis is basic understanding and, hopefully, melioration of a structure or a structure-generating process. Melioration and, more ambitious, optimization assumes that a goal can be defined. This almost trivial prerequisite, taken for granted in most cases, does in fact require further consideration if we leave the usual textbook examples. We will show several practically important cases where the definition of the 'correct' target function is by no means stringent. There are even systems to which such a thing as a target function in the common understanding cannot be attributed in principle.

2.5.1 A reflection on optimality

What is the optimal society, or an optimal public transportation system? Even if we leave out any subjectivity-related optimality definition, the problematic nature of this question becomes clear by the similar one: What is an optimally adapted plant or animal?

If we ponder on those questions we will notice some properties of systems that make them hard to treat in any optimizing framework. The predominant problem with the mentioned systems is their multitude of principally sensible quality measures that all seem more or less arbitrary, both by definition and relative weighting. If we look at biology, the concept of an optimally adapted animal has been ever-changing over time, space and context, as the phylogenetic tree [WA] and the diversity of the living world tell us quite impressively. Nevertheless there seems to be a kind of directional evolution in the development of a certain species along its historical course. This observation in turn tempts into assuming some kind of teleological positivism [Wikc] by interpreting these developments as directed towards a final state — implying that a targeted melioration or even optimization is feasible. There are objective observations, however, that suggest this assumption not to be realistic.

If we compare different answers of biological evolution to very similar environmental conditions in spatially far away corners of the world we find quite disparate concepts that obviously have proven as well fitting: The rather large number of different marsupial¹ species on the Australian archipelago inhabit habitats very similar to those presently dominated by ordinary (placental) mammals in the rest of the world. There is an interesting *a posteriori* interpretation of the differing developmental traits of those concepts, though [Wikb]: 'The early birth of marsupials removes the developing young much sooner than in placental mammals, and marsupials did not need to develop a complex placenta to protect the young from its mother's immune system. Early birth places the tiny new-born marsupial at greater risk, but significantly reduces the risks associated with pregnancy, as there is no need to carry a large fetus to full-term in bad seasons. Because a newborn marsupial must climb up to its mother's nipples, the otherwise minimally developed newborn has front limbs that are much better developed than the rest of its body. This requirement is responsible for the more limited range of locomotory adaptations in marsupials than placentals; marsupials must retain a grasping forepaw and cannot develop it into a hoof, wing, or flipper as some groups of placental mammals have done.'

¹Marsupials are mammals that fawn premature offspring and feed it during the first part of its further development in a skin pouch outside the womb. The young remain there for a period corresponding to the late stages of fetal development of a placental mammal. The best-known representatives of these creatures are the Australian kangaroos.

It may be estimated that many systems that a human optimizer would like to meliorate fall into a similar category: Interpreting a historical development in terms of an identified melioration effect succeeds quite well, but a prediction of further development, or even the derivation of an optimality point is not viable. Many ways of potential development will lead to similar qualitative improvement with respect to singled-out target functions, but with unforeseeable side-effects on further 'total quality defining' target aspects.

Even if a multitude of potentially important targets can be defined in an optimization effort, they may change their relative weights or even appear or disappear during the process of ongoing melioration, when effects neglected at the start of the consideration become decisive in later stages, or limiting boundary conditions eventually dissolve. Let us inspect the century-long development process of a typical technical system — the melioration of power plants.

When the vast expansion of electrical power usage started in the earlier decades of the last century the foremost quality of plants was the cycle efficiency while the total power output of a given design and its adaptability to changing load situations were considered to be secondary aspects. Due to still lacking large-scale transportation facilities usually near-by primary energy sources, as e.g. local coal mines, played an important role. If exhausts were regarded at all it was mostly limited to the pollution by particulates in the nearest neighborhood of the power plant, usually resulting in quite high chimneys to enforce better mixing of exhausts with clean air. With growing size of the energy market and individual request amounts the cost efficiency became more and more important as the electrical energy entered a steadily growing range of end consumer products (like household aids, electrically operated industrial machines etc.). The growing demand created dependencies on far-away suppliers of primary energy sources that led eventually to the first world-wide "oil crisis" in the seventies of the last century — a new competing goal of power station (or at least: power provision) valuation sprang into existence: the provision security.

A relief from this situation was sought by the alternative construction of nuclear energy plants, lessening the dependency of the industrial countries on fossile energy imports. Only regarded from a technical point of view this emerging technology promised large amounts of cheap electrical energy. But in various countries the opposition to nuclear power production was under-estimated². The additional expenditures for shielding and securing the required infrastructure against opposing groups, as well as fulfilling politically imposed requirements of strongly raised security assessments, made the price of electrical energy production explode relative to earlier calculations. Had the social and political side effects been foreseen they might have been counteracted on a definitely lower impact level. But as history tells this chance had been missed.

As a short-term effect the return of the fossile energy burning power plants, meanwhile very cleanly operating with respect to particulate emissions, collided with the rather recently emerging concerns of volatile exhausts leading to the emergence of a strong renewable energy movement. While on the long run this method of consumable energy production on the basis of renewables will be the only one proving practicable, even on a mid-range perspective it might lead to raising costs in the immediate future, potentially leading to undesired side effects on economic and social development.

Summarizing this historical development of striving for the best (i.e. cheapest, most productive, stable-operating) power plants we obtain some insights into the emerging optimization process peculiarities. The persecution of a central main objective for certain periods leads to the upcoming of additional goals, not having been regarded before, that grow to play limiting and counteracting roles in the whole process. So it is unpredictable in detail, changing in an abrupt manner, and thereby has the properties of dynamical structure generation. Even if later-on appearing goals had

²This description is mainly abstracted from the historical development in Germany. Other countries, like France, exhibited significantly different developments.

been taken into account at the start of the long-term development they would not have contributed to a more pointed and less crisis-driven development, as the additional targets would have led to *an inferior outcome at an earlier point in time*. Alternatively proposed solutions, anticipating later objectives, would not have been competitive at earlier points in time. This, finally, leads to the conclusion that the system 'effective electrical energy preparation' is one without a objectively defineable long term target.

There are many more systems and optimization objectives that are comparable in this respect. Here, just one other shall be discussed that is surely familiar to (almost) any reader: the development of easy and efficient individual locomotion over larger distances.

Earth-bound individual locomotion is strongly coupled to the improvement of passenger cars. Due to the technical success of the respective automobile development and the present exorbitant use of them (at least compared to the times of early development) traffic jams, too few parking space, fuel costs, and air pollution became severe limiting factors, only counteracted by additional enforced technical and logistic developments. Those changing boundary conditions in turn modify the development paths modern automotive technology is taking. The structure generation process of individual mobility thus assumes the properties of a dynamical structure. While in earlier times maximum velocity, spacyness and comfort were anticipated future development goals, public attention and customer focus has strongly shifted towards economical operation, crash stability, unobtrusiveness and endurance, with an upcoming perspective on navigation support and automatized, instantaneous avoidance of traffic jam conditions.

While this description primarily holds for passenger cars, an almost opposite direction of development is observed for the two-wheeled versions of automotive devices, the motor bikes. In earlier times they were mainly seen and used as a simple means of individual transportation, with mostly cost-effectiveness, solidity and ease of maintenance being the predominant goals. Contemporary aspects strongly deviate from this view, bringing individuality, enforced demonstration of power etc. to front. Extrapolating views of the past, a typical recent motor bike would have been quite probably attributed as too heavy, too difficult to service, and too expensive in former times.

We may interpret these views in terms of static vs. dynamical structure generation and sustenance. We may regard the goals of a technical system, eventually in combination with an (averaged) subjective view of a potential evaluator, as static as long as we deliberately blind out longer-range perspectives. A power station or a car *may be* optimized with respect of a certain goal, or even a set of goals, if its or their apprehension is regarded as constant. Changes of apprehension frequently manifest themselves in their distribution within a society, quite often with a gradient on individual age, leading to a certain type of 'generation conflicts'. So, in an unrestricted view, neither sociological nor technical structures lend themselves to an analysis of optimality. However, they do so in a restricted view. In a limited time horizon as well as in a limited location technical structures can be optimized. We shall discuss the constraints of the optimality algorithms in the next section.

2.5.2 Constraints in optimum seeking

In a partial view, many structures and structure generating processes lend themselves to melioration and even optimization endeavors. Their actual executions are determined by a number of constraints, also frequently referred to as boundary conditions.

Parameter ranges as boundary conditions

Preparing a system for (numerical) optimization, we first model the causal dependencies of one or more target function values upon a set of configuration parameters. These dependencies need

not necessarily be direct, sharp, or simple. In most practical cases complex multiple, non-linear and unsharp dependencies are quite common. Configuration parameters are typically *bounded*, i.e. they cannot assume arbitrary values: Modeling the uplift of an airship we have to lower-bound the weight of the lift gas to that of helium, as there is no one lighter and at the same time safer gas than that. Modeling the performance of a gas turbine power station (depending on the properties of its operative parts, like turbines, pumps, heaters etc.) we have to limit the maximum temperatures of the underlying thermodynamical process to the maximum values that the material contact faces of the structural elements can tolerate.

These constraints usually limit the attainable values of the target qualities as well. In many cases it is not obvious, though, what extremal target function values can be obtained. Here the underlying model must be queried, either by direct mathematical evaluation of dependencies, or by numerous samplings of target values depending on respective sets of input variable settings.

The notion of limitations or bounds in the model definition, implying attainable target values, seems rather trivial at first sight. As soon as the potential ranges of configuration variables depend on each other, this grows into a sometimes very complex problem, though. Revisiting the power station example, we may observe that the maximum surface temperature on the gas turbine blades and thus the efficiency depends on the amount of additional cooling gas input through fine air ducts inside the blades. Additional cooling gas reduces the maximum output of the turbine again. Furtheron the available cross section for such air ducts is limited by the required material stability for the very fast spinning blades, thus limiting even more design variables. While this description is by no means exhaustive it should suffice to demonstrate the complex dependencies of target function value on the limiting configuration factors. This makes it rather difficult to estimate the influence of them on available configuration parameter spans. It directly brings us to the problem of potentially excluding interesting regions of (unbounded) configuration parameter space without the ability to grasp it.

Bounds are not necessarily maximum or minimum limiting conditions. They may as well appear in the interior of an otherwise already limited configuration parameter range. As an example, let us consider another thermodynamic system: H_2O under environmental pressure at its boiling point. In that state we will either observe it in the rather dense form of liquid water (density approx. 960 kg/m^3) or in its diluted state of water vapor (density 0.6 kg/m^3). There is no in-between, and system configurations requiring a homogeneous form of water with an intermediate value cannot be realized. If one tries to approach such a state the formerly homogeneous material (either water or vapor) separates into two phases with strikingly different mechanical and thermodynamical properties. This process is highly important for appropriate layouts for turbomachinery, compressors and even tubing systems.

Superstructures

As soon as structural variants can be defined for the realization of a system we enter the domain of discrete optimization. For the matter of argument we restrict our discussion here to a problem class containing integral as well as continuous configuration variables. A typical representative of such a problem is the chemical engineering process optimization shown in section 2.4.3, page 29. Besides several continuous variables there are structural alternatives, each of them defining their individual subset of dependent parameters.

A frequently adopted method to manage the space of potential structural alternatives is the definition of a so-called superstructure³. It contains all allowed structural variants as binary

³The concept of superstructures is much more encompassing than the pure interconnections of structural alternatives if we use it in a more transdisciplinary context, but for the reason of argument we only understand its meaning in the

switching options, expressed in terms of allowed parameter values of 0 and 1 respectively. As with continuous variables, simple decision variants are readily intelligible. But as soon as intermingled and mutually dependent switch settings are required the system may become highly complicated. Again, we note the problem of ascertaining that every desired parameter setting, here with respect to structural alternatives, should be able to be reached during the optimization process. If this is not the case we may unknowingly exclude important ranges of the “true” (i.e. pragmatically feasible) configuration space from systematic optimization.

In general, such superstructures occur in combination with continuous variables, thus creating mixed-integer linear or non-linear optimization problems. Especially the latter are known for their algorithmic complexity with respect to rigorous and affirmed location of their respective global optima. They continue to be subject of current research and development.

One of the major problems in such systems is the fact that configuration parameters are related to each other in a hierarchic way: The continuous parameters defining the properties of a chemical engineering plant in some switchable sub-tree (like pressures, temperatures, sizing, material selections) only *exist* if the sub-tree is switched on. Otherwise they are completely meaningless. In simulative practice, usually the complete decision space with all potentially needed continuous parameters is set up. When a respective branch of the flowsheet is switched off they are just considered existing, but meaningless with respect to the target function value(s).

Open-ended structures

A boundary condition of a system’s structure can also be defined implicitly and self-adaptingly. In that case the system’s higher-level semantic definition must automatically be interpreted and translated into a functional model amenable for computational evaluation. As this is a highly complex task there are not too many reported examples for this methodology. Therefore we would like to mention this kind of approach more as a desirable future goal in system modeling than a broadly applied technique.

As an example we may re-inspect the chemical engineering flowsheet as shown in fig. 2.11 on page 30. The conventional method to model such a system is the superstructure approach discussed in the last paragraph. Here connections of pre-determined apparatuses are fixated, each being defined by its respective set of design variables. Mathematical algorithms may then be applied to identify a best possible setting of the variables, leading to the respective value(s) of target function(s). An alternative semantic definition of the flowsheet may be given by just describing the effect that is desired from a concatenation of various kinds of generalized unit operations. We might say: “The lower-boiling fraction separated in ‘Flash 1’ is to be fed into a distillation column (representing $y_4 \rightarrow 1$)” and so forth, of course favourably in a more formalized language structure, but with the same level of (im)preciseness. Layout changes can be expressed relatively easy in such an abstracting language, as on this level of generalization it is, per definitionem, not necessary to define precise configuration parameter settings in details. Instead we rely on computational intelligence for re-adapting them at least to a certain measure of quality. If we restrict ourselves to such system definitions there is ample room for modifying the layout of the interconnection structure, just by changing some elements of our description. We may introduce additional items or take items or whole branches out.

Compared to the more mathematically oriented superstructure approach it is somewhat simpler to assure the definability of all sensible *structural* variants. But as the optimizing effort needs the assistance of computerized hierarchically underlying continuous parameter definition intelligence, we in turn cannot be sure that for all structural variants the complete scope of continuous parameter

narrower sense usually attributed in the area of chemical engineering.

settings is realized, so the general problem of potentially excluding interesting configuration space ranges remains.

It is very difficult to subject this kind of semantics-based structural redefinition to a rigorous, algorithmic system optimization. In the rather few documented cases this approach was put to work evolutionary concepts were applied, e.g. in [Sci96] for the optimization of a chemical plant. A kind of intermediate method is the application of 5th generation programming languages, like Prolog [Wik06b]: Here the human programmer can concentrate on the logical structure of the simulation model, without having to bother about the inner looping and variables control. These inner structures are organized by the Prolog interpreter automatically as soon as a simulation is run. But for realistically complex simulation models this claim of complete procedural automation has to be significantly limited, almost withdrawn, as the high-level Prolog program has no notion of practical relevance or additional, mostly unsharp contextual limitations that would impose respective evidence limits on the governing resolution methodology of 'back-tracking'. These issues have rather to be dealt with by manually introducing so-called 'cuts' that manifest lower-level reasoning based on individually occurring parameter settings.

2.5.3 The chore of choice: Trading off conflicting goals

The problem of multiobjective optimization is not new. Already in 1871 Vilfredo Pareto investigated in his work the problems of multi objectiveness and suggested the mathematical concept of best compromise solutions[Par71]. Many, if not almost all, melioration problems of the real world exhibit more than one target function. Cars are not only valued for their maximum speed but also for a low fuel consumption value, a high reliability, low maintenance costs etc.. A delivery service not only needs to deliver the goods in time but also is expected to handle them carefully, to be competitive in its costs, or to obey time windows for pick-up and delivery as close as possible. Usually such sets of target aspects cannot be satisfied optimally for each criterion at the same time. Let us, in this respect, reconsider the passenger car example: Satisfying the desire for high maximum speed will require large and complex engines, as well as elaborated braking equipment, which in turn leads to higher investment and maintenance costs. A part of the raised costs may be caught by applying less solid constructive layouts but will in turn lead to lower reliability. Almost any scenario of constructive change will stress different aspects of the arbitration effort of „optimizing everything at the same time“.

Accordingly, the systematic, objective optimization of a system will eventually reach a bound it cannot transgress, and the domain of the subjective choice sets in: If two or more conflicting goals cannot be resolved at the same time, the individual valuation is the last resort to choose *the one* solution suggestion that is to be realized in actual utilization. Nevertheless there is a *set of best solutions*, frequently called the „Pareto Set“ of the given problem (fig. 2.14) in honour to Vilfredo Pareto. It represents those solution suggestions where one target aspect can be improved only by impairing at least one other. Mathematically, it creates a mapping

$$\vec{x} \in D \implies \vec{f}(\vec{x})$$

with \vec{x} representing the vector of system design variables and $\vec{f}(\vec{x})$ the vector of dependent individual target function values. This mapping can be discrete (in the case that only distinct configuration parameter settings are feasible) or continuous (if there are arbitrarily adjustable configuration parameters like temperatures or pressures in a technical system). If there is a large number of potential distinct settings the individual frontier points may be represented by an idealizing continuous boundary representation that may be handier for consecuting evaluation purposes. These pareto sets exhibit some properties that shall be discussed further.




Figure 2.14: Graphical representation of a two-dimensional Pareto front. Both the actually determined target function combinations (starred positions) and an idealized, continuous front line are given for the non-dominated solution propositions.

If we begin with scrutinizing a single Pareto front for a technical system with continuous variables, we frequently observe a qualitative target arbitration behaviour schematically shown in fig.2.14 for two targets. As an example we may take a residential heating system, with the inlet temperature into the hot water piping as the governing variable to trade heating costs, f_{costs} , vs. lagging of temperature adaptation in the heated rooms, f_{lag} . Both targets should be minimized. Raising the inlet temperature will yield a faster response of the system, i.e. f_{lag} will decrease. But at the same time energy losses rise and hereby produce higher costs, i.e. f_{costs} will increase, and vice versa. The dependency of both target functions on the temperature is, due to the underlying thermodynamic laws, smooth but conflicting.

As a certain setting must be chosen, at least for a given point in time, it is the question how the arbitration is individually performed. A classic method is the linear superposition of target function values, defined by a combined quality function Q_{eff} (see fig. 2.15):

$$Q_{eff} = \frac{\alpha f_{costs} + \beta f_{lag}}{\alpha + \beta} .$$

Any such aggregation will define a linear superposition of the originally separate functions that may be expressed as a tilted new “optimization axis”, equivalently expressed by arbitration iso-lines perpendicular to that axis. The angle of such iso-lines therefore may as well define the individual arbitration model. If, as an extreme example, almost only heating costs are considered, the arbitration axis would almost coincide with the abscissa of our plot, and iso-lines would be near to parallel to the ordinate. Every single pair of system responses in f_{costs} and f_{lag} , resulting from a respective setting of the piping inlet temperature, is projected onto the new optimization axis, leading eventually to the identification of an ‘optimal point’. Changing the subjective focus and hence the relative weighting of the individual goals will result in a shifted optimum and, together with it, in a different temperature setting for the heating system.




Figure 2.15: Classical arbitration of two target functions by linear superposition. The general concept is transferable to more than two targets. Explanation see text.

If there are discrete configuration parameters in a pragmatically defined optimization problem, up to the point that strongly differing solution concepts compete with each other, the combined pareto front tends to become more complex and exhibits additional features. Let us investigate a rather simple but illustrative, idealized example: If we have to decide which means of transportation to use for traveling from a location A to a location B, we are interested in minimizing two target functions: the duration of the trip, represented by the inverse average velocity, and the costs of the trip. It is evident that taking a higher velocity transportation method will usually induce raised costs, as the investment, fuel, and maintenance costs of an airplane, for example, exceed that of a bicycle, or the repetitive consultation of a shoemaker's services.

Mainly depending on the total distance to be traveled, quite different vehicles can be chosen for a certain trip. Only for very small distances the plane is no alternative, and only for the really long distance trips bicycle riding and walking is ruled out. For each vehicle an individual pareto set may be expected: By applying a high technological input (like low friction bearings, aerodynamical spokes, streamlined frames etc.), a respective bicycle can be made substantially faster compared to cheap ones from the superstore. But even for such high-tech bikes the ultimate limiting condition is the power input by the human, defining the maximum velocity even for highest-priced items. Similar arguments hold for the other vehicles as well, thus creating individual pareto sets of their own, expected to penetrate each other at specific, but case-dependent, front intersection points.

The resulting pareto set, shown in figure 2.16 as the purple-marked border line, exhibits some interesting properties. First, as a superposition of the individual pareto front lines, it is not necessarily differentiable, as the frontiers usually intersect with non-zero angles. Second, it does not even need to be steady, as adjacent partial pareto fronts need not touch each other. In our qualitative example, the pareto curve for the car starts already at higher velocities than a bicycle rider will ever reach. So, if at the position of that unsteadiness a small increase of trip costs is accepted there is a leap in transportation velocity by the switch to the other method of transportation. Such




Figure 2.16: Creation of a structured, qualitative Pareto set of two target functions by superposition of individual sets of structurally different solution approaches. Explanation see text.

leaps can happen in any direction.

If we apply our method of linear superimposing target functions to this plot (as shown in fig. 2.15) we will, in this very restricted quality assessment view, arrive at the conclusion to use either plane or bicycle. Tilting the superposition assessment axis from one extreme to the other will leave out any other means of transportation. On the other hand this implies that rather extreme values of costs and velocity are regarded as equally sufficing at the switching point from one solution alternative to the other. In a practical point of view that is obviously not what we intended to obtain. The solution key is the acceptance of *non-linear* assessment functions: In practice we honor the *simultaneous* fulfillment of target function qualities, as long as it does not lead to a major deterioration of the linearly superimposed solutions. Mathematically this may be represented by defining the combined quality function Q_{eff} with an additional, non-linear term

$$Q_{eff} = \frac{\alpha f_{costs} + \beta f_{lag}}{\alpha + \beta} + \gamma \cdot f_{costs} \cdot f_{lag} \quad (\gamma < 0 \text{ for minimization}),$$

With an increasing non-linearity factor γ we add stress to the simultaneity of advantageous individual target function settings. It is almost impossible, though, to express such subjective, nonlinear preference settings in an objectivized plot. Rather, it is sensible to report the complete (quantative) Pareto set calculated for an optimization problem to an individual or a group of human decision takers. On that basis every involved person can decide how far he or she would compromise on a back-laying, but simultaneously near-advantageous configuration of target values.

Apart from regions of intermingled subsystem Pareto fronts, most neighboring Pareto set elements originate from neighboring configuration parameter settings. Small changes in target function values are usually caused by equivalently small changes in the system parameter settings, as is obvious in our residential heating example.

If a number of n target functions is considered the dimension of the „pareto front“ is $(n - 1)$. Therefore, presenting pareto sets for two conflicting targets is not too elaborate: One can — in case of discrete system layouts — just either sketch the possible realizable target function combinations for the non-dominated solution suggestions, or present an abstracted 'front line' of those solutions for densely populated sets (fig. 2.14).

Addressing more than two targets is more difficult: Even if there still is a pseudo three-dimensional representation for three targets, its visualization and interpretation is not so simple. A fast display, freely rotatable in every direction by computer mouse interactions, can help to get a correct notion of the non-dominated front face items for the given problem. Even if an interaction of the decision maker with a computer is necessary for this, the display of the pareto front is basically static. Besides view angle changes no active work of the computer is required.

Proceeding to four and more targets changes this situation significantly: There is no static possibility to display a 'pareto front' with so many dimensions. How is a decision maker then able to adhere to the concept of pareto optimization; how can he get a notion on the pareto elements of a given problem, to be able to trade off the goals according to his or her individual preferences, taking into account the limits of feasibility?

Here the decision maker support requires active work on the computer's side. As the pareto set cannot be displayed as a whole, individual interactions must help investigating it. A first approach might be to set all but two of the targets to fixed values temporarily and see what (partial) pareto front shows up for the two unrestricted values. The drawback of this approach is evident: For n criteria there are $\binom{n}{m} = n!/2! \cdot (n - 2)!$ possibilities of combining them. So four criteria create six two-dimensional views, five criteria ten, and six criteria already 15. In addition this method only creates 'cuts' perpendicular to the variables' axes, so interesting combinations of criteria results are not considered.

An interesting alternative to those static display efforts are interactive spider plots or star plots as discussed in detail in the case studies 3.2 and 3.3. Here the identified pareto-optimal solutions of a given problem are displayed together with controls for each individual target function range. Operating these controls for one target immediately shows the consequences for the attainable ranges of the other ones, thus helping the human investigator to obtain 'a feeling' on the arbitration space that he or she can choose the favourite solution from.

In its last resort, multi-criteriality always is a matter of subjective choice: If two or more conflicting goals cannot be resolved at the same time, with every objective criterion being already considered, the individual taste is the only guide to choose *the one* solution suggestion that is to be realized.

Another possibility of identifying relative decision makers' preferences is the preemptive request on several (assumed, not necessarily existing) combinations of target values. In choosing between pairs of suggested properties relative weightings may be derived. This is no way to perform the choice itself, though, and the real-life selection may be influenced by other objectives than isolated pair-wise decisions.

2.5.4 Objective and subjective quality

Up to this point only objective target functions have been discussed, being historically predominant in technical optimization contexts. But we should not close our eyes to the fact that there are many problems with intrinsically subjective goal functions. Most of them involve human perception and senses, and subjective preferences play a crucial role. Examples are the production of a coffee mixture with a special taste and smell, the various sounds produced by a vehicle, the development of perfumes, or digestibles with intended smells or tastes.

For illustration, a taste optimization is discussed briefly, based on [Her97]. How can a coffee mixture with optimal taste be created? As an appropriate methodology an evolution strategy has been identified, working on subjective assessments by comparative taste attributions of testers. There are no objective criteria which can be quantified, but coffee testers nevertheless have the task to produce a coffee with an intended taste. This presupposes a sense of taste which is trained in long times of practice. Additionally, the evolutionary procedure does not need a metric scale, instead it only relies on a *relative* comparative attribution of preferences. The only thing the test person has to do is to compare tastes pair-wise.

An evolutionary strategy for optimal taste determination may be set up by the following steps: Five randomly mixed coffee mixtures are given to the coffee testers. The mixtures ratios are not known by the testers. The task is to find out the coffee mixture coming closest to an offered reference mixture. The mixture of the coffee approaching the reference taste best is taken as a basis for five new mixture variations. This looped procedure converges after a few follow-ups — the testers cannot differentiate anymore between the quite similarly tasting test mixtures, although there are objectively differing mixing ratios, also with some distinct differences to the goal mixture. The relative assessments become ambiguous. Nevertheless the intended goal is reached: The desired taste has been approximated with the limits of subjective difference perception.

From this example we may abstract some characteristic conditions for optimization procedures with subjective goal functions:

- Optimization goals usually possess a considerable fuzziness.
- An optimization goal is crucially influenced by the group of assessing persons and their subjective notion of quality fulfillment.
- Subjectively shaped goal functions are usually multi-criteria functions (see below).
- Objectivity, which is often asked for and aimed at, can only be realized in a limited way.

An important aspect of procedural development in treating subjectivity-biased melioration processes is the identification of more or less objectively determinable goal definitions that might replace the present subjective ones. If this succeeds even at least partially, some additional objectivity may be gained and the range of undecidedness be reduced somewhat. Some illustrating examples of presently subjectivity-dominated issues may be taken from the field of technical acoustics:

- Acoustic quality of larger rooms, like concert halls, auditoria, or classrooms,
- adjustment of hearing aids for hear-impaired individuals,
- walking noise emission of modern laminate floors,
- acoustic quality of loudspeakers, musical instruments, bells, etc.,

- sound design as well as sound quality, used as acoustic visiting-card, for special technical products, like automobile doors.

Some subjectively perceived sound qualities are relatable to objective measurable quantities, although the overall perception of sound as an acoustic quality is very complex in nature. There are influential subjective factors as individual conceptions, hearing habits, etc. Consequently, sound perception cannot simply be described by objectively measurable sound pressure levels (or loudness), pitch (or timbre), and reverberation in a room. They correlate in some cases quite well, but in other cases there are striking discrepancies. As of today's knowledge objective assessments can replace subjective assessments only in exceptional cases. However, subjectively superior acoustic qualities often coincide with certain settings of objectively measurable properties.

Two facets of subjectivity play important roles in this respect. First, human beings do not perceive sound events in complex situations as a separable input. Instead, acoustic perceptions are embedded into non-acoustic environmental influences: The acoustic perceptions in an concert hall will also be dependent on optical perceptions, on architectural aesthetic effects, the expectations in acoustic performance, the comfort of the taken position (climatic conditions, disturbing influences, seat comfort), and on the psychosocial condition of the listener (missing acceptance of a musical piece, personal physical condition). Second, the assessment is influenced by the group of test listeners and by its type. The assessment changes when performed by a group of laymen or experts, trained or untrained listeners, active or passive participants (listeners to music or musicians).

Insights from the field of acoustics may be transferred to other subjectivity-influenced domains, like textual relevance attribution or climatic comfort issues. If we try, for example, to identify similar texts in a larger collection of writings, the *practical* attribution of similarity is strongly dependent on the relevance attribution of included passages as set by the assessor. Even if larger parts of two compared texts diverge, individual sections with subjectively decisive keywords found in both may lead to the conclusion that they are closely related. In the area of text comparison there is a chance of objectification by applying a frequency-of-incidence measure on individual words. But as ultima ratio only the valuating subject can judge the similarity.

The climatic comfort constitutes a more objective, although not sharply defined subjective target function, e.g. in the assessment of buildings. While the individual perception of coziness may differ substantially in a given room there are at least well-tested statistical comfort models, taking into account different dimensions like environment temperature, radiation temperature imbalances of surrounding walls, environmental humidity, airflow around the human body, and some other influencing factors. Even though the resulting comfort model does not represent the perception of a certain single individual in a room, it very well describes the probability density distribution of the comfort attribution of a larger number of people. So, if a building is not designed for a certain small number of nameable persons, but to serve an initially not known number of anonymous inhabitants like office workers, these quasi-objective quality measures may well serve the needs of building designers in their effort of trading off different target qualities like costs, comfort, flexibility and such.

2.6 Mathematical Methods

Model analysis requires mathematical methods. Particularly challenging tasks and the solution of multicriterial optimization problems, of problems with uncertainty and of problems carrying other completed features of the real world. We here differentiate between rigorous mathematics and evolutionary algorithms.

2.6.1 Rigorous mathematics

2.6.2 Challenges for Contemporary Evolutionary Algorithms

General Remarks

Does one need more than one optimization method? Or, stated differently, is there an optimal optimization method? Following from the No Free Lunch theorem (NFL, Wolpert and Macready [WM97b]), in the general case — without clearly specified task — there is not. For every single task, creating a specialized method would be advantageous. Unfortunately, this requires (i) a lot of effort, and (ii) extensive knowledge about the treated problem, and is thus not practicable. Alternatively, two strategies are usually followed when tackling a ‘new’ optimization problem:

- Adapt an existing algorithm to the problem in its current form, and/or
- model/formulate the problem appropriately for an existing algorithm.

Whereas ‘traditional’ mathematical optimization approaches mostly favor the second approach, it may provoke unwanted side-effects: One has to make sure that the most important features of the original problem are taken over into the model. E.g., matching the problem to an existing algorithm may obscure its real global or good local optimizers so that they become unreachable for the optimization algorithm. Besides, many existing algorithms require the problem to fulfill properties it obviously or possibly does not do, e.g. continuity and differentiability. Particularly, in cases where computing the quality value of a solution candidate requires running a complex simulation software, one seldomly knows in advance which properties the underlying (unknown) objective function possesses.

When nothing more than quality determining response values for any set of input variables are known for a problem, we speak of *black box* optimization. In the single-objective case, the common notion of an objective function and its global optimum/global optimizers — as given in eqn. 2.1 for unconstrained problems — is still useful. However, global optimizers, i.e. the set of input vectors \mathbf{x} for which $f(\mathbf{x})$ is optimal, cannot be determined analytically. An empirical trial and error method is the only way to find them.

$$f^{*G} = \min\{f(\mathbf{x})|\mathbf{x} \in X\} \quad (2.1)$$

The black box concept immediately leads to *direct search* methods, i.e. methods that only utilize objective function responses and “do not ‘in its heart’ develop an approximate gradient”, as Wright [Wri95] puts it. As far back as in the 1960s, many direct search methods have been invented, e.g. the famous *Nelder-Mead simplex algorithm* [NM65]. At the same time, the first steps into the world of *evolutionary computation* (EC) were taken, presenting very simple versions of what is now subsumed under the unified denotation *evolutionary algorithms* (EA). These do not only use bio-inspired heuristics, they also employ randomness. However, the extensive use of random numbers and the fragmentary theory supporting EAs may be considered a drawback. Nevertheless, these optimization methods have demonstrated their problem solving capability in numerous real-world applications.

Interestingly, in recent years, the mathematical optimization community has again shown increased interest in direct search methods, e.g. Kolda et al. [KLT03]. This may have to do with (i) the fact that these techniques simply did not go extinct on the practitioners side, and (ii) improved theoretical analysis methods that now help tackling heuristic algorithms. In computer science, the growing field of *randomized algorithms* is exclusively dealing with algorithms employing random numbers — not only in optimization. Motwani and Raghavan [MR95] give an overview.

This section targets at introducing the main EA concepts and specialized techniques for three important application areas: Multiobjective optimization, optimization under uncertainty, and multimodal optimization. These are relevant to the topic of this book as they are closely interrelated and often encountered conjoined in real-world applications.

Historical roots

Although there have been precursors in proposing the utilization of evolutionary concepts for optimization tasks, as e.g. Bremermann [Bre62] (also see Fogel's fossil record [Fog98]), invention and development of the first evolutionary algorithms is nowadays attributed to a handful of pioneers who independently suggested three different approaches.

- Fogel, Owens, and Walsh introduced evolutionary programming (EP) [FOW65], at first focused at evolving finite automata, later on modified into a numerical optimization method.
- Genetic algorithms (GAs), as laid out by Holland [Hol73], mainly dealt with combinatorial problems and consequentially started with binary strings, inspired by the genetic code found in natural life.
- Evolution strategies (ESs) as brought up by Rechenberg [Rec73] and Schwefel [Sch75] began with solving experimental engineering problems by hand using discrete/integer parameters, but turning to real-valued representations when numerical problems had to be solved.

In the early 1990s, a fourth branch of evolutionary algorithms emerged, explicitly performing optimization of programs: Genetic programming (GP), suggested by Koza [Koz92]. Since about the same time, these four techniques are collectively referred to as evolutionary algorithms, building the core of the evolutionary computation (EC) field.

What is an evolutionary algorithm?

Today, there is little doubt about components and general structure of an EA. It is understood as population based direct search algorithm with stochastic elements that in some sense mimics the organic evolution.

Besides initialization and termination as necessary constituents of every algorithm, EAs consist of three important factors: A number of search operators, an imposed control flow (fig. 2.17), and a representation that maps adequate variables to implementable solution candidates.

Although different EAs may put different emphasis on the search operators mutation and recombination, their general effects are not in question. Mutation means neighborhood based movement in search space that includes the exploration of the 'outer space' currently not covered by a population, whereas recombination rearranges existing information and so focuses on the 'inner space'. Selection is meant to introduce a bias towards better fitness values; GAs do so by regulating the crossover via mating selection, ESs utilize the environmental selection.

A concrete EA may contain specific mutation, recombination, or selection operators, or call them only with a certain probability, but the control flow is usually left unchanged. Each of the consecutive cycles is termed a *generation*. Concerning the representation, it should be noted that most empiric studies are based on canonical forms as binary strings or real-valued vectors, whereas many real-world applications require specialized, problem dependent ones.

For an in-depth coverage on the defining components of an EA and their connection to natural evolution, see Eiben and Schoenauer [ES02], and Eiben and Smith [ES03], and Bäck, Fogel, and Michalewicz (eds.) [BFM97].

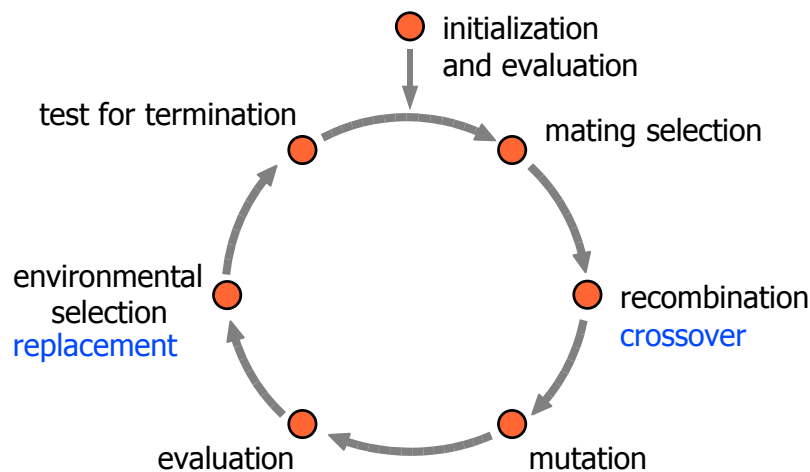


Figure 2.17: The evolutionary cycle, basic working scheme of all EAs. Terms common for describing evolution strategies are used, alternative (GA) terms are added below.




Figure 2.18: The evolutionary cycle of a two-membered (1+1) evolution strategy.

Evolution strategies

In the following, we introduce the most important canonical ES variants for single objective optimization, which serve as basis for more specialized algorithms later on.

The first ES, the so-called (1 + 1)-ES or two membered evolution strategy, uses one parent and one offspring only. Two rules are applied to these candidate solutions:

1. Apply small, random changes to all variables simultaneously.
2. If the offspring solution is not worse (in terms of its function value) than the parent, take it as the new parent, otherwise retain the parent.

Schwefel [Sch95a] describes this algorithm as “the minimal concept for an imitation of organic evolution.” The (1 + 1)-ES (fig. 2.18) is applied by many optimization practitioners to their optimization problem and included in our analysis for three reasons: (i) It is easy to implement, (ii) it requires only few exogenous parameters, and (iii) it defines a standard for comparisons.

The first $(1 + 1)$ -ES used binomially distributed mutations for integer variables (Schwefel [Sch65]). These have been replaced by Gaussian mutations for continuous variables. Rechenberg [Rec71] already proposed a simple rule to control the mutation strength, the so-called 1/5 success rule. Note that even this simple ES requires the specification of at least four factors, namely the adaptation interval, the required success rate, the step size adjustment factor⁴, and the step size starting value.

Population based ESs will be presented in the following paragraph. They use μ parents and λ offspring. Note that whereas GAs rely upon a start population uniformly scattered in a closed search region, ESs — even if population based — may be started around any start vector like standard optimization algorithms, without lower and upper bounds for the variables.

Variation in ESs The use of populations enables an extension of the rather simple 1/5 success rule to control the mutation strength (Schwefel [Sch75]). Beyer and Schwefel [BS02b] propose some guidelines derived from the philosophy of Darwinian evolution to design these variation operators.

1. A *state* comprises a set of object and strategy parameter values $(x^{(t)}, s^{(t)})$. *Reachability* demands that any state can be reached within a finite number of iterations. This feature is necessary to prove (theoretically) global convergence.
2. Variation operators (mutation and recombination) should not introduce any bias, e.g. by considering only good candidate solutions. Variation operators are designed to *explore* the search space in contrast to selection operators that exploit the gathered information. Recombination works, according to Beyer [Bey95], mainly as gene repair operator, not only as building block collection mechanism.
3. *Scalability* is the third criterion that should be fulfilled by variation operators: Small changes of the representation should cause small changes in the function values.

The standard ES recombination operators produce one offspring from a family of ρ parent individuals (usually $\rho = 2$). Consider a set of μ parental vectors of length N , representing either object or strategy parameters:

$$\{(x_{11}, \dots, x_{1N}), (x_{21}, \dots, x_{2N}), \dots, (x_{\mu 1}, \dots, x_{\mu N})\}. \quad (2.2)$$

Two recombination schemes are commonly used in ESs. Both use a set $\mathcal{R} = \{r_1, r_2, \dots, r_\rho\}$, that represents the indices of the mating partners. It is constructed by randomly (uniformly) choosing ρ numbers (with replacement or not) from the set $\{1, 2, \dots, \mu\}$. *Discrete recombination* selects the entries of the offspring randomly from \mathcal{R} , whereas *intermediary recombination* averages the ρ corresponding values of all mating pool members in each component of the newly generated vector.

Mutation is applied to the recombined intermediate solution. Mutation in multimembered ESs is a self-adaptive process that relies on the individual coupling of endogenous strategy parameters with object parameters. After being varied as described above, the strategy parameters (standard deviations, also called mean step sizes or mutation strengths) are applied to mutate the object parameters. To illustrate this procedure, algorithms with one common σ are considered first. To prevent negative standard deviations, mutation of this σ should be done multiplicatively. Beyer and Schwefel [BS02b] discuss an additional argument for a multiplicative mutation of the mutation

⁴This is a constant factor c with $1 \leq c \leq 0.85$, the lower bound being theoretically near-optimal for simple model problems like the sphere model.

strength on the sphere model. It can be shown, that in expectation σ should be changed by a factor that only depends on N . Therefore, the mutation operator can be implemented as

$$\sigma^{(t+1)} = \sigma^{(t)} \cdot \exp(\tau z), \quad (2.3)$$

where z is a realization of an $\mathcal{N}(0, 1)$ distributed random variable. The parameter τ is the so-called *learning rate*. The object variables are mutated next:

$$x^{(t+1)} = x^{(t)} + w, \quad (2.4)$$

where w is a realization of an $\mathcal{N}(0, \sigma^{(t+1)})$ distributed random variable. The multiplicative mutation scheme for one σ can be extended to several strategy parameters $\sigma = (\sigma_1, \dots, \sigma_N)$. Schwefel [Sch77] proposes the following extended log-normal rule:

$$\sigma^{(t+1)} = \left(\sigma_1^{(t)} \exp(\tau z_1), \dots, \sigma_d^{(t)} \exp(\tau z_N) \right), \quad (2.5)$$

where z_i are realizations of N standard normally distributed random variables, $i \in 1 \leq N$. This mutation scheme employs a single learning rate τ for all strategy parameters. An alternative procedure that utilizes a global and a local learning parameter τ_0 and τ , respectively, is suggested by Bäck and Schwefel [BS92]. Self-adaptive correlated mutations have already been introduced in 1974, see Schwefel [Sch81] and Schwefel [Sch87].

Selection in ESs Selection should direct the evolutionary search toward promising regions. In ESs, only candidate solutions with good function values are allowed to reproduce. The replacement (environmental selection) process is deterministic in contrast to the random processes used in GAs. This selection scheme is known as *truncation* or *breeding selection* in biology. The κ -selection scheme takes the age of candidate solutions into account: Only candidate solutions that are younger than κ generations may survive, regardless of their fitness. For $\kappa = 1$ this selection method is referred to as *comma-selection*: only offspring individuals can reproduce. The κ -selection is referred to as *plus-selection* for $\kappa = \infty$: Both the offspring and the parents belong to the mating pool. The plus-selection is an elitist selection scheme, because it guarantees the survival of the best individual found so far.

Rechenberg introduced the first multimembered ES, the so-called $(\mu + 1)$ -ES. It uses μ parents and one offspring and is referred to as the *steady-state* ES. Schwefel introduced the $(\mu + \lambda)$ -ES, in which $\lambda \geq 1$ candidate solutions are created each generation, and the best μ out of all $\mu + \lambda$ individuals survive, and the (μ, λ) -ES, in which the parents are forgotten and only the best μ out of λ candidate solutions survive. A birth surplus is necessary for the (μ, λ) -ES, that is $\lambda > \mu$. Schwefel et al. [SRB95] and Beyer and Schwefel [BS02b] provide a comprehensive introduction to evolution strategies.

Designs of Evolutionary Algorithms Table 2.1 summarizes important ES parameters. The reader is referred to Bartz-Beielstein [BB03] for a more detailed description.

Ways to Cope with Uncertainty

In many experimental optimization problems, quality values can only be obtained with a limited degree of exactness. Falsely calibrated measurement instruments, inexact scales, scale reading errors, etc. are typical sources for measurement errors. If the quality values are the output of stochastic simulations, then the measurements may be exact, but some of the models' variables

Table 2.1: Default settings of exogenous parameters of a “standard” ES. Source: Bäck [Bäc96]. Problems may occur, when these “standards” are blindly adopted and not adjusted to the specific optimization problem.

Symbol	Parameter	Range
μ	Number of parent individuals	\mathbb{N}
$\nu = \lambda/\mu$	Offspring-parent ratio	\mathbb{R}_+
$\sigma_i^{(0)}$	Initial standard deviations	\mathbb{R}_+
n_σ	Number of standard deviations. N denotes the problem dimension	$\{1, N\}$
τ_0, τ	Multiplier for mutation parameters	\mathbb{R}_+
ρ	Mixing number	$\{1, \mu\}$
r_x	Recombination operator for object variables	{intermediary, discrete}
r_σ	Recombination operator for strategy variables	{intermediary, discrete}
κ	Maximum age	\mathbb{R}_+

are random variables. Thus, there may be at least two sources of uncertainty: Noise on quality values or noise on input variables (especially the ones that cannot be controlled by the optimization process).

Consider a sequence of traffic signals along a certain route or elevators’ movements in high-rise buildings. *Optimization via simulation* subsumes all problems in which the performance of the system is determined by running a computer simulation. If the result of a simulation run is a random variable, we cannot optimize the actual value of the simulation output, or a singular performance of the system. Instead, the average performance of the system in question may be a practicably sensible goal of optimization via simulation.

The stochastic output in optimization via simulation complicates the selection process for direct search methods. The efficiency of the evaluation and selection method is a crucial point, since averaging over repeated runs reduces the efficiency of the optimization process. In the context of selection and decision making, the terms “candidate” and “point” will be used synonymously. A well established context where these error probabilities are analyzed is hypothesis testing.

Noise makes it difficult to compare different solutions and select the better ones. Noise affects the selection process in evolutionary algorithms: In every iteration, the best μ out of λ candidate solutions have to be determined.

Wrong decisions can cause *stagnation* of the search process: Over-valuated candidates—solutions that are only seemingly better—build a barrier around the optimum and prevent convergence. Or, even worse, the search process can be *misguided*: The selection of seemingly good candidates moves the search away from the optimum. This phenomenon occurs if the noise level is high and the probability of a correct selection is very small.

One may attempt to reduce the effect of noise explicitly. The simplest way to do so is to sample a solution’s function value n times, and use the average as estimate for the true expected function value. This reduces the standard deviation of the noise by a factor of \sqrt{n} , while increasing the running time by a factor of n . Further means used by evolutionary algorithms to cope with noise are averaging techniques based on statistical tests, local regression methods for function value estimation, or methods to vary the population size [Sta98, Bey00, SK00, Arn01, BSS01, BBM04].

Following Bartz-Beielstein [BB06], we present a taxonomy of elementary selection methods to tackle with noise. Depending on a priori knowledge, selection schemes can be classified according to the following criteria:

Threshold: subset selection – indifference zone.

Termination: single stage – multi stage (sequential).

Sample size: open procedures – closed procedures.

Variances: known – unknown, equal – unequal.

The goal of subset selection is the identification of a subset containing the best candidate. It is related to screening procedures. *Subset selection* is used when analyzing results, whereas the *indifference zone* (IZ) approach is used when designing experiments. The sample size is known in subset selection approaches, it is determined prior to the experiments in the indifference zone approaches. *Single stage* procedures can be distinguished from *multi stage* procedures. The terms “multi stage” and “sequential” will be used synonymously. The latter can use *elimination*: If inferior solutions are detected, they are eliminated immediately. Selection procedures are *closed*, if prior to experimentation an upper bound is placed on the number of observations to be taken from each candidate. Otherwise, they are *open*. Furthermore, it is important to know whether the variance is common or known. Bartz-Beielstein [BB06] discussed similarities and differences of these approaches. He also analyzed threshold-based procedures, which were successfully applied to noisy, dynamic functions, e.g., in elevator group control. Threshold rejection increases the chance of rejecting a worse candidate at the expense of accepting a good candidate. It might be adequate if there is a very small probability of generating a good candidate.

Multiple Objectives

As was noted above for many problems of high practical interest in science and engineering, several possibly contradicting objectives shall be pursued simultaneously. In science and social life we are confronted with many examples, such as in chemical process engineering, where the productivity of chemical reactors is in contrast to their loss during the start up and shut down phases. In the textile industry, a similar conflict arises for the production of fabrics. Figure 2.19 shows a simple discrete example. Total elongation (F_1) and extensibility (F_2) of the fabric shall be improved, by means of maximizing F_1 and minimizing F_2 . All objectives are sufficiently defined and in this case pointwise quantifiable. Their values are determined by three adjustable control factors (decision variables): Number of knitting skewers (x_1), number of knitting rows (x_2) and number of weft threads (x_3) per inch. The challenge for a multi-objective optimization algorithm consists of finding decision variable value sets that fulfill all objectives as well as possible.

In this context, the Pareto [Par96] concept of optimality proved fruitful, as discussed in many general terms above. During the beginning of an optimization run, it is often not hard to find solutions that simultaneously improve both objectives. However, if an objective can be improved further only by worsening an other objective, a solution is called *Pareto-optimal*. Due to different possible preferences concerning the single objectives, this leads to a set of Pareto-optimal solutions, each of them representing a valid optimal solution for the multi-objective problem (MOP). Figure 2.19 shows six solutions in the decision variable space (a) and the objective space (b) for the fabric improvement example. In this example, the decision variable space is discrete and constrained as indicated by the surrounding solid line. Consequently, there is only a finite number of possible objective value combinations. Direct comparison of solutions 5 and 6 shows that the former improves on F_1 without changing F_2 . According to the Pareto dominance concept, solution 5 *dominates* solution 6. However, pairwise comparison of solutions 1 to 5 does not result in recognizing any such domination as improvement in one objective always comes along with worsening in the other. These solutions are therefore indifferent to each other, hence incomparable or *non-dominated*. If due to problem-specific constraints no further improvements can be obtained (solutions 1-5 are on the border of the feasible region) the set of all non-dominated solutions

represents the *Pareto Set* in the decision space and the *Pareto Front* in the objective space. Since in each case only one solution can be realized, preference information of a decision maker (DM) must be used next to select the final solution of the MOP.

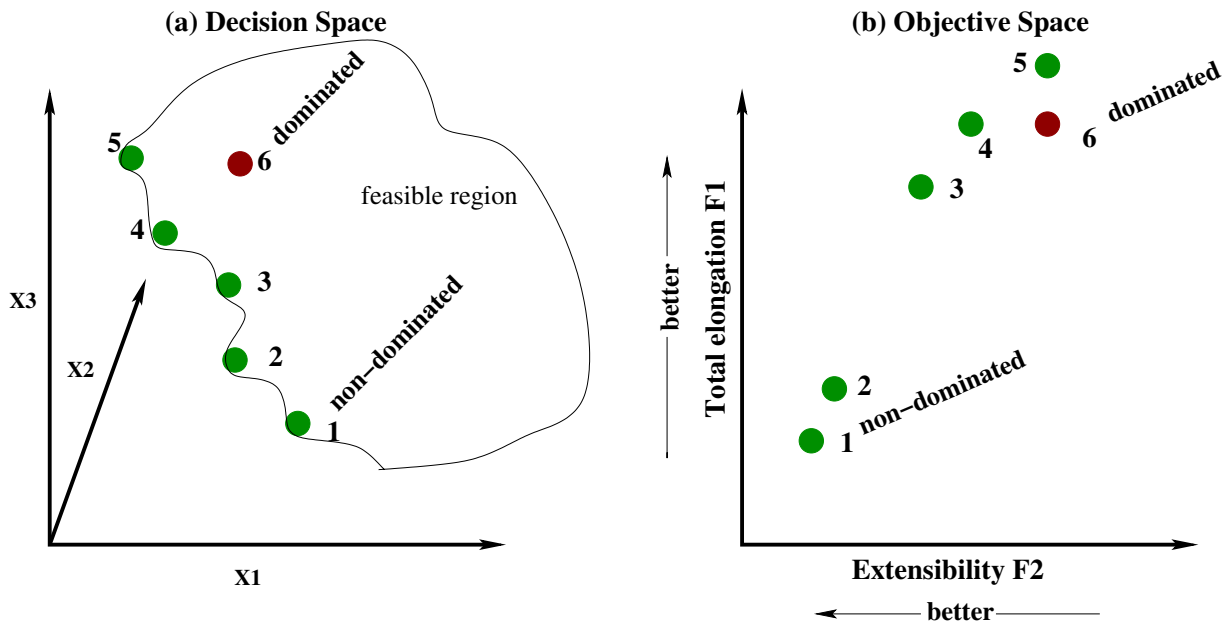


Figure 2.19: The Pareto-dominance concept. (a) Decision space, (b) objective space

Why Use Evolutionary Algorithms? Problems with several conflicting criteria have been treated for many years, e.g. with a considerable variety of techniques developed in Operational Research. Concise overviews of existing approaches can be found in Achilles et al. [AEN79] and Miettinen [Mie98]. Usually one tries to reduce the MOP into a single-objective problem, so that it can be solved by means of methods from single-objective optimization. One possible approach consists of choosing a single criterion as main objective, and transform the other objectives to constraints with lower or upper bounds. Without specific knowledge of the problem, the choice of concrete upper and lower bounds suffers from arbitrariness. Alternatively, one may try aggregation-based approaches. These combine all criteria into a single, parametrized one. The aggregation can be accomplished by any combination of arithmetical operations (i.e. a weighted sum), according to some understanding of the problem. However, these techniques have several limitations. Some of them are e.g. susceptible to the shape (convex/concave) of the Pareto front, others to its continuity (connected/disconnected). In addition, most of the ‘conventional’ approaches are only able to compute one single non-dominated solution per run. Searching for a representative set of non-dominated solutions requires a restart with different external parameter settings and different starting points for each run.

Evolutionary algorithms are robust search methods, whose success and failure is by far less susceptible to the shape or the continuity of the Pareto front. Their greatest advantage is that they are able to provide a point-wise approximation of the whole Pareto front in one go by employing cooperative search of a whole population.

Algorithm Design If one regards the development of the *evolutionary multi-objective (EMO) algorithms* within the last two decades, then the rise of suggested approaches is impressing. The largest well-known collection of existing approaches was arranged by Coello Coello and contains over 1900 entries [Coe06]. A common classification of all EMO-algorithms comes from

Masud [HM79]. Depending on the time at which the preference information from the DM is used, four classes can be differentiated: (i) Non-preference, (ii) a-priori, (iii) interactive, and (iv) a-posteriori. In the following, this classification is not discussed in detail as most EMO-algorithms can be assigned to the last category. The optimization process takes place before any preference information is incorporated. This entails a clear task definition: Find a representative set of non-dominated solutions as close (convergence) as possible to the *Pareto optimal set/front*. Additionally, the resulting approximation has to exhibit a good distribution of solutions in terms of both spread and uniformity - usually described by the term of *diversity*. The aim of this section is to give an overview of the main methods that have been developed in order to achieve these goals.

Fitness Assignment When moving from single-objective to multi-objective optimization while applying EAs, the most important changes to be made concern the selection operator and especially the fitness assignment. In EAs, the fittest individuals have better chances to survive and reproduce. For single-objective optimization, only one scalar fitness value exists. However, in the multi-objective case we have to deal with a fitness vector. Since EAs need a scalar to work on, generally two design decisions must be made: On the one hand, this vector must be scaled to enable for EA selection, and on the other hand the two conflicting tasks of convergence and diversity shall be respected. But how to assign the fitness of an individual in order to express suitability towards both goals? We can roughly divide the existing answers into two categories:

Combined Fitness Assignment: Fitness is assigned such that the fitness value represents convergence and diversity at the same time.

Single Fitness Assignment: Fitness assignment respects only one goal. Usually, this is convergence, as in the single-objective case.

Aggregation-, performance-, and Pareto-based approaches belong to the first category. Aggregation-based approaches are the most traditional as well as simplest possibility. Recently, performance-based fitness assignment strategies are successfully used to evaluate the fitness of a new individual in relation to the entire population. For example, the S-metric selection (SMS)-EMOA utilizes the well-known S-metric (hypervolume) to calculate the fitness of an individual. This measure is commonly used to evaluate the performance of an EMOA. It respects proximity to the Pareto front as well as diversity of the solution set.

Pareto-based approaches use the Pareto dominance concept itself for fitness assignment. Differences between these approaches arise in the methods employed to exploit the partial order. According to Zitzler et al. [ZLB03], this kind of information can be divided into: (i) *Dominance rank*: The number of solutions in the population that dominate the solution under consideration, (ii) *dominance count*: The number of solutions in the population that are dominated by the solution under consideration, and (iii) *dominance depth*: The rank of the solution in the non-dominated sorted population. The latter approach is utilized by many successful algorithms, e.g. the *Non-dominated Sorting Genetic Algorithm (NSGA)-II* by Deb and others [DAPM00b]. Dominance rank was first employed by Fonseca and Fleming in their *Pareto envelope-based algorithm (PESA)* [FF96]. Today, a multiplicity of methods are based on this principle, see for example Bosman and Thierens [BT05]. Dominance depth and dominance rank are successfully combined in the *Strength Pareto Evolutionary Algorithm 2 (SPEA2)* approach by Zitzler and others [ZLT01b].

However, most of these algorithms apply a *secondary fitness assignment strategy* that serves the goal of diversity. In most cases they try to incorporate density information into the selection process (mating/environmental), according to the rule: The smaller the density of individuals

within a neighborhood, the larger the chance of an individual to reproduce. Figure 2.20 shows the three most frequently used methods: *Kernel-based*, *grid-based* and *nearest-neighborhood* measures. Fitness sharing, as e.g. used in NSGA, is a kernel-based strategy. The distance of an individual to all other individuals in the population is calculated and summed up. These values are then used to deflect the evolutionary search out of densely populated regions. Grid-based techniques as e.g. utilized by the *Pareto Archived Evolution strategy (PAES)* of Knowles and Corne [KC99], employ hypergrids to define neighborhoods within the objective space. The more individuals in a box, the heavier they are penalized (see fig. 2.20). Nearest neighborhood techniques as used in SPEA2 and its variants calculate the distance between an individual and its nearest neighbor in order to estimate the neighborhood density.




Figure 2.20: Most common diversity preservation strategies in EMOA.

Criterion-based approaches represent the second category of fitness assignment strategies. They all share the same basic idea: The fitness value of an individual is determined by only one of the criteria according to the goal of convergence. However, the choice of a single criterion for any individual shall be reconsidered repeatedly (in each generation). As thereby parts of the population are selected according to different criteria, it is hoped that the goal of diversity can be achieved indirectly (see Schaffer [Sch84] and Laumanns and others [LRS98]).

Representations and Variation Operators Design and analysis of representations and corresponding genetic operators is prevalent in the field of evolutionary computation. Often, an adept combination of all components determines the system's success or failure. This insight is ubiquitous in the case of single-objective optimization. However, in multi-objective optimization, the conceptual approaches are still mainly concerned with the selection operator. Research focusing on variation operators or representations remains rare. Some recent approaches are: Rudolph [Rud98] and Hanne [Han99] who investigate control mechanisms for the mutation strength in the multi-objective case. Grimme and Schmitt [GS06] focus on recombination operators that produce diverse offspring in each generation.

Elitism Elitism preserves previously attained good solutions from one generation to the next. The prime example of an elitist algorithm in the single-objective case is the 'plus'-selection ES. In the

multi-objective case two types of elitism are used: Maintaining elitism in the current population, as is already done in the single-objective case, or doing so in an archive (secondary population) that stores non-dominated solutions externally. Archive contents may or may not be integrated again into the optimization process (Zitzler and others [ZLB03]). Of vital importance is the criterion used to control replacement of archive members, the most commonly used of which is the dominance criterion. It leads to an archive of non-dominated solutions, relative to all solutions generated during a run.

Future Perspectives As has been hinted to in the previous paragraphs, a lot of work remains to be done on EMOAs. We briefly discuss the currently most promising paths:

Investigating representations and variation operators: Büche and others [BMK03] show that the interaction between selection and search operators is often not co-ordinated well, and that approximation of the Pareto front cannot be done with arbitrary precision. Further on, there is the dilemma of stagnation with good diversity of the solution set on the one hand, or arbitrarily exact approximation of a few points on the Pareto front. We conjecture that this trade-off between convergence and diversity can be attributed to the fact that variation operators cannot simply be taken over from the single-objective case and that changing only the selection operator is not sufficient to meet the requirements of multi-objective optimization.

Focusing on the region of interest (ROI): In the last years, most EMO researchers focus on algorithms that are able to find the whole Pareto front. However, in practice, the decision maker is only interested in a specific region of the Pareto-front. Focusing on a region derived from user preferences may help to increase convergence speed and/or quality and also simplify solution selection by the DM later on.

Parallelism: Considering the suitability of EAs working in a parallel manner, one should expect that the development of parallel approaches stands only at the beginning. Apart from first successful attempts to convert the state-of-the-art algorithms into a parallel version [OHMW02], an increasing number of parallel approaches has been published only recently [CC05, MMSK04].

Parameter tuning: Attaining good parameter settings for a given problem-algorithm combination currently is one of the hot topics in single-objective optimization [BB06]. It is necessary to adapt those techniques for the multi-objective case in order to avoid the commonly used manual parameter tuning and provide important insight into parameter interactions.

Multimodal Problems

Although, during the last decades, many empirical and most of the theoretical studies in EC have been devoted to simple test problems with only one extremal point, the great majority of practical applications requires optimization in far more complex fitness landscapes. Multimodality—the presence of more than one locally optimal point—requires a shift from a hill-climbing oriented towards a global perspective. At the top of the hill, the need arises to somehow ‘escape’ the associated local optimum. This may be done in two different ways. Either, one tries to save as much positional and learned (step sizes/mutation strengths) information as possible and, preserving this information, attempts to jump over the neighboring valleys. Or, one completely gives up the current search space location and performs random initialization again. For mutation strengths getting larger and larger, the former scenario more and more resembles the latter.

However, if the treated optimization problem is not available in a closed algebraic form, detecting the arrival at a local optimum may not be trivial, depending on the employed variable representation. Combinatorial and binary encoded optimization problems come with a natural minimal step definition which enables enumeration of the neighborhood. For real-valued representations, eqn. 2.6 specifies a necessary and sufficient condition for a local optimum, with \mathbf{x}^{*L} meaning its search space location, $d(\mathbf{x}, \mathbf{y})$ a distance metric, and ϵ the maximal distance to tested neighboring search points. Nevertheless, the bounded but still infinite neighborhood cannot be completely explored efficiently and one has to rely on the strong causality assumption (Rechenberg[Rec89]: similar causes entail similar effects) to identify local optima at least in probability.

$$\mathbf{x}^{*L} \text{ is local minimizer iff } \exists \epsilon : \forall \mathbf{x} \in X : d(\mathbf{x}, \mathbf{x}^{*L}) < \epsilon \Rightarrow f(\mathbf{x}^{*L}) \leq f(\mathbf{x}) \quad (2.6)$$

Strongly related to the notion of local optima is the one of basins of attraction; these encompass the search space portion leading to an optimum if the steepest descent is followed. For this local search process, efficient approximation methods are known, e.g. quasi-Newton algorithms. However, identification of different basins is even more difficult than local optimum detection if no further information regarding size and/or location of the basins is available. The key property of multimodal optimization methods is thus how efficient they are in finding the different search space regions that contain the best local optima.

Canonical population based EAs perform global and local search at the same time, gradually narrowing their focus to the most promising regions, and more sooner than later to a single basin of attraction (e.g. Preuss, Schönemann and Emmerich [PSE05]). From the discussion above, it becomes clear that the ability to explore multiple promising regions—either concurrently or sequentially—is decisive for obtaining well performing EA variants. But for a given limit of available computational time, these always have to face the global vs. local search tradeoff like any other global optimization algorithm.

One possible way to speedup local optimization, so that more effort can be diverted to search space exploration, is to hybridize EAs with existing local search methods. These approaches are subsumed under the term *memetic algorithms* (MA) that was introduced by Moscato [Mos89]. A recent overview is given by Krasnogor and Smith [KS05], together with a suggested taxonomy.

Most other specialized EAs strive for enhanced global search capabilities by means of at least one of the following three techniques:

Restarts are utilized to enhance the chance of reaching the/a basin of attraction of the global optimum. As an example, an efficient restart CMA-ES for multimodal problems has been suggested by Auger and Hansen [AH05]. Multistart methods obtain potential solutions consecutively, and every new instantiation may be provided with search results of completed previous runs. They avoid the problem of jumping into a neighboring good region by giving up the current search space location completely.

Diversity maintenance aims for a uniform distribution of individuals over the whole search space. Comparing relative or absolute distances of solution candidates and applying clustering methods are common means to prevent overlapping search paths and promote good search space coverage. Diversity may be held up explicitly or implicitly. Following Eiben and Smith [ES03], explicit means that active measures are taken to model the distribution of search points in the desired way, whereas implicit stands for deliberately slowing down information exchange by restricting recombination or selection/replacement. Classical island models provide implicit diversity maintenance by building relatively independent subpopulations. Spatially structured EAs [Tom05] do so by restricting the effect of recombination and

selection operators to the local neighborhood. *Shifting balance* GAs by Oppacher and Wineberg [OW99] exemplify explicit diversity maintenance as they prevent subpopulation overlap which is measured by absolute population distances.

Niching methods also strive for a suitable spread of search points, only on the level of basins of attraction. As Mahfoud [Mah95] points out, it is the aim of *niching* algorithms to detect separate basins and keep them in focus of the search. Unfortunately, *basin identification* within an EA is not easy and prone to error, so that endogenously retrieved basin information is highly unreliable and nonexistent when the optimization starts. Crowding by De Jong [De 75] and fitness sharing by Goldberg and Richardson [GR87] are regarded as the classical niching methods. The former employ relative, the latter absolute distances. These have been carried further e.g. by Li et al. [LBPC02], Streichert et al. [SSUZ03], and Shir [Shi05], but still the radii employed for detecting search points located together in a basin remain problematic. Only few approaches integrate fitness topology information into the basin identification process, e.g. the *universal evolutionary global optimizer* (UEGO) by Jelasity [Jel98], Ursem's multinational GA [Urs99], and the sample-based crowding method proposed by Ando et al. [ASK05].

It shall be noted that solving multimodal problems is related to tackling constrained or multiobjective ones. Removing constraints from a problem by transforming it by means of (metric) penalty functions (see e.g. Michalewicz and Schoenauer [MS96] and Coello Coello [Coe02]) as commonly done in EC most often leads to multimodal problems even if the original problem was unimodal.

In multi-objective optimization, the focus has been mainly on the objective space for a long time. Today, it becomes increasingly clear that population movement in the decision (search) space heavily depends on the multimodal search properties of the applied optimization algorithms (Preuss, Naujoks and Rudolph [PNR06]).

Conclusions

May it be (or not) that one day there is no more need to invent new optimization tools because we have got the best tailored ones already for every possible real-world problem. May it be (or not) that then the dream of hardliners has come true that all of these best tailored methods can abstain from using pseudo random numbers for deciding upon the next iteration in the search for the solution. But, contemporary tools are still well advised not to rely on deterministic algorithms alone. That is, why an idea from the early days of digital computers is still alive, i.e., the idea to mimic procedures found in nature that obviously have led to remarkably effective systems or subsystems. One may think that nature had enough time to achieve a good solution by means of pure chance, but time has always been scarce when there are competitors, and the way nature finds its way is much more sophisticated.

Anyway, it is a matter of fact that evolutionary algorithms have become widely used in practice since their invention in the 1960s and even found their way into articles in the field of theoretical computer science. Their domain of application are 'black box' situations, where the analysis of the situation at hand does not help or is too costly or dangerous, i.e., in case of experimental design and even computer simulation of nonlinear dynamic systems and processes.

Chapter 3

Case studies

The study of structure generating processes is an interdisciplinary subject. Case studies from the various fields reveal similarities as well as differences in the methods of solution and also in the goals of analysis. The following selection serves to give an overview over the treat of applications. It is by no means exhaustive. The individual articles are written by experts in the various fields with a strong emphasis on illustrating the points elaborated in the preceding chapter on the transdisciplinary foundations. Readers wishing to enhance their understanding and to proceed more thoroughly to the specific disciplinary details are put in the position to do so by consulting the cited literature.

3.1 Production Engineering: Optimal Structures of Injection Molding Tools

Jörn Mehnert, Thomas Michelitsch, Klaus Weinert

Molding is an important and frequently applied production technique of mass production: Most plastic components from micro switches to bumpers of cars are manufactured via injection molding. Metallic structures such as aluminum gear boxes or stairs of escalators are produced via die casting. All these products are generated by injecting a hot material into a die. In zinc die casting the injected metal has a pouring temperature of about 435 °C and the pouring temperature of aluminum varies from 620 °C (thick walled) to 730 °C (thin walled) [Has06]. Aluminum and zinc are examples of the most prominent metallic casting materials. The material cools down, solidifies, and the workpiece is ejected from the die before a new process cycle can start. In this process the layout of mold temperature control designs is decisive for efficient die casting. The efficient cooling of a die is one of the main factors to reduce the cycle time and, therefore, to increase the cost effectiveness of a mass production tool. Improvements of up to 50 % [Wie03] are possible, because today the layout of the bores carrying the cooling liquid in casting tools is typically done manually. The layout design is based on expert knowledge and rules of thumb. Experts have a good estimate about where to position cooling bores properly. Heuristic rules are quite common in practice. These rules are fast and work very efficiently. Unfortunately, only little of this knowledge is available in the form of explicit mathematical rules [Zö97]. Solutions for complex problems may not be satisfying as existing heuristic rules deduced for simpler tasks may contradict each other. If problems with blowholes or insufficient surface qualities arise as a consequence of an inferior bore layout the die requires frequently very expensive and time-consuming remanufacturing. Due to the complex creation process die casting tools cost up to five hundred thousand Euro. Therefore, well designed cooling circuits help to significantly reduce manufacturing costs of the die, minimize cycle times, increase tool life and improve the workpiece quality.




Figure 3.1: Example of a cooling bore design in a mold with two independent cooling circuits. The die surface is shown golden-red, while the mold itself is only represented by the bores contained in it.

Systematic optimization requires quantitative measures for the quality of a proposed solution. Therefore, the experts knowledge together with physical properties have to be modeled, leading to the exact definition of quality criteria, the choice for efficient problem spaces, data structures,

and efficient quality criteria evaluation methods. Due to the multiobjective character and the high complexity of the problem, powerful optimization algorithms have been used. Very useful techniques for solving difficult problems are evolutionary algorithms.

These algorithms are able to find surprisingly good solutions by clever analysis of the quality functions. Sometimes even experts are surprised about the new solutions found and get motivated to use new approaches. There are several advantages of evolutionary algorithms. The algorithms are rather robust against changes in the fitness functions.

Many classic optimization algorithms – this is also true for older evolutionary algorithms – are only able to solve single objective problems. Although often the single objective case is complex enough, many real world problems have a multiobjective character. Problems may have contradictory aims, i.e. one objective value cannot be improved without deteriorating an other. The reduction of all criteria to only one value is possible only in special cases. In the case of uncertainty about the preferences of a user toward one or an other criterion, a set of best compromise solutions should be presented, as already discussed in section 2.5.3.

Multiobjective evolutionary algorithms can be used either with a posteriori or with a priori techniques. A posteriori approaches follow the idea of uncertainty about a perfect solution. Their goal is to generate best compromise solutions, so called Pareto optimal solution sets, from which an expert can choose according to non-formalized ad hoc knowledge. In case of well defined multiobjective goals sometimes also scalar evolutionary algorithms are useful.

Several multiobjective evolutionary algorithms are mentioned in literature. A very good evolutionary scalar optimizer is the Evolution Strategy (ES) [Sch95b]. It can be used for multiobjective problems technique via aggregation. State of the art a posteriori multiobjective optimizers are the NSGA II [DAPM00a, DAPM00b] and the SPEA 2 [ZLT01a]. All three techniques have been tested on the mold temperature design problem.

Another important issue, when evolutionary algorithms are applied to real-world applications, is computational speed. Stochastic algorithms tend to need a lot of fitness function evaluations. Especially for difficult problems with high dimensions, restrictions, a difficult multimodal character etc. lead to high population sizes or long runs of the stochastic search algorithm. Therefore, an efficient design of the evaluation functions is necessary. The definition of fast evaluable quality functions is particularly necessary in the case of mold temperature control design. Surrogate functions or a well motivated simplification of the problem can lead to a significant speed up without much loss in quality.

Finally a systematic search for good parameter adjustments of the multiobjective search methods is necessary to improve the search process itself and its results. The best algorithmic parameter settings for various problems often differ from the standard. In order to compare multiobjective optimization algorithms that generate Pareto fronts, it is common to introduce scalarization methods, so called metrics. The application of statistical design methods helps to systematically, but experimentally, meta-optimize the parameter settings, leading to robust, statistically sound and optimized parameter adjustments. Here, existing classical and new statistical approaches have been applied to multiobjective as well as to single criterion optimization of the mold temperature design problem.

Modeling Aspects

Geometric Aspects

Molding tools are composed of various parts such as sliders, bores, pillars, the die surface etc. The cooling or heating of a certain area of a die is realized by introducing channels into the tool that lead water or oil into the vicinity of the die surface. These channels form circuits. The number of

bores and circuits depend on the complexity of the molding tool. Simple structures need few bores and a single circuit only. Complex tools may contain more than thirty bores and more than five circuits. The circuits can be operated independently from each other and can be used to heat up or cool down the tool. Heating up is important for reducing thermal stresses particularly for large tools in die casting. Dynamically changing strong thermal stresses can lead to increased tool wear or even to tool breakage. In the following, cooling and heating circuits are subsumed under the term cooling circuits because for the present discussion only the geometric structure of the bores is important.

The number of the bores should not be too high because manufacturing of deep hole drilling bores is cost and time intensive. Deep hole drilling machines for large molding tools show restrictions with respect to the orientation of the bores. Often only horizontal orientations of the bores are possible while the die can be rotated on a table freely around the z-axis. Although some modern deep hole drilling machines do not have these restrictions, a design tool for realistic bores should be able to obey them.

Another restriction in the bore directions is related to machining problems of intersecting bores. Bores that intersect each other in small angles cannot be manufactured safely by deep hole drilling. Therefore, it is preferable to design bores that intersect each other in large angles — ideally orthogonal.

The number of bores should also be kept small because any bore implies machining risks and costs and also weakens the molding tool. A certain wall thickness around each bore is necessary to reduce the probability of deformation or even breakage of the tool. In the design of large die casting tools rules of thumb are fast and well established but also very case depending and cannot be transferred linearly for small tools because the physics of the temperature follows nonlinear laws.

The simplest structure of a cooling circuit is a consecutive sequence of bores. This is also the most prominent structure. More complex cooling strategies such as the branching of bores or so called finger bores for local cooling are more complex but lead to a higher efficiency of the cooling. Shortcuts with other cooling circuits are not allowed. A grouping of inlets and outlets of several circuits is also desired from a practical point of view.

Dead ends of bores used in circuits with more than two segments are sealed with plugs. Dead ends appear because each bore of a circuit has to be machined by deep hole drilling from outside. Plugging is necessary to seal a circuit from leakage. Plugs can be placed in any part of a bore. After setting a plug it seals a hole permanently. In some cases a plug can also be used to seal a die surface that has been penetrated by a bore. In general, drilling through the die surface is not allowed and should be avoided whenever possible because the subsequent machining, e.g. grinding, is very cost intensive.

Typically, dies and molds are designed with CAD systems. There exist various techniques for describing the geometric elements of a tool in the CAD system. The most general way is the triangulation of the surfaces. CAD-surfaces can be exported from CAD systems, e.g. in standard STL or IGES format and can be imported into a simulation software. The consecutively used simulation software then utilizes triangulation for describing the complete molding tool. The bores are modeled as virtual cylinders. Actually, it is possible to use arbitrary shaped triangulations of a die's surfaces. For simplification reasons the tools are generally assumed to be rectangular blocks. The cooling cycles are modeled as sequences of consecutive cylinders. In the mathematical abstraction these sequences are simple polylines. Each cooling circuit can be defined uniquely by the position of the vertices of the polyline. The start and end points of the polyline always lie on the border surface of the molding tool. The model contains the complete bores for the testing of collisions and efficient machining. Plugs are defined by their position at the polyline.

The drilling direction for the machining of a bore is not unique. Therefore, always both possible drilling directions can be modeled and analyzed.

Quality Criteria and Evaluation

The cooling of an injection mold has to satisfy several aspects [Meh05]. The most important idea is to generate cooling circuits that keep the cycle times in the later usage in mass production as low as possible while the machining of the circuits should be as cost efficient as possible. Already these two basic aspects are contradictory. On the one hand the amount of bores enhances the cooling effect and hence reduces the cycle time. On the other hand the costs increase with the number of bores. The efficiency increases with the number of bores but the costs do also. Hence, the basic problem is multiobjective with conflicting criteria. The efficiency of the cooling of the circuits can be increased by changing the layout of the cooling bores. The corresponding design problem also has a multiobjective character because the cooling efficiency of each bore should be both uniform and strong. A uniform cooling distribution along the die surface is important e.g. to reduce internal stresses within the workpiece. An intense cooling is necessary for short cycle times. The relative distance and orientation of a bore to a point on the die surface is decisive for the cooling efficiency of the bore. Larger distances generate more homogenous temperature distributions on the die surface but impede the absolute cooling intensity and vice versa. The length of a bore also has a certain influence on the cooling effect. The longer the bore the better the cooling effect but the higher the costs for the manufacturing of the bore.

The complete expert knowledge for modeling the layout of cooling circuits can hardly be captured by an algorithm. Nevertheless, a lot of aspects of the current 'fuzzy knowledge' can be formalized and made available for mathematical use. Many values can be deduced from physical or economical properties of the machining process. Some knowledge can be formalized as quality criteria, e.g. average temperature of the die, others as restrictions, e.g. possible bore directions of the drilling machine. Basic properties such as the bore directions may also be encoded directly in the data structures. For example if the bores would have been restricted to lie in a single plane the problem would have been encoded different from the arbitrary 3D case as in the model used here.

The definition of the fitness functions can also benefit from efficient encodings. Here, the meaning of efficient encoding refers to fast fitness function evaluation as well as to a reduced difficulty for the optimization strategy. A low dimension of the search space and a continuous fitness landscape with only few restrictions is preferable.

For a technical realization of the efficient description of the casting expert's knowledge, physical properties of the thermal flow in metal, economical assumptions about the manufacturing process as well as mathematical heuristics have been used. A realistic assessment about a sufficient level of detail is a helpful technique to reduce computation times. These approximations can be used to choose parameters of the geometric model of the tool such as the number and size of the triangles used to describe the die surface. Approximations have also been applied to obtain a very fast estimation of the cooling effect of the bores.

The quality of the cooling induced by a bore is determined by the local intensity of the cooling effect to the die surface. Additionally, a globally uniform heat extraction is important to get e.g. workpieces with a minimum of internal stress and, therefore, high qualities. Global and local cooling are conflicting goals. Additional conflicting goals such as the number and length of bores have to be taken into account. The layout of the bores has to fulfill technical and machining restrictions. The bores must not intersect neither with the die surface nor with any other bore. Of course, all bores have to be within the tool geometry. This restriction seems trivial in concept, but it is a not at all so with respect to the simulation system boundary control implementation.

Additional machining restrictions such as limited bore direction angles have been addressed in the optimization tool.

The evaluation of the cooling effect can be calculated very exactly by Finite Element Methods (FEM). Unfortunately, this approach needs long calculation times. FEM has the advantage that it calculates the complete temperature distributions within the tool for all times. It is necessary to use high numbers of nodes, because a low resolution can lead to relevant discretization errors when the bores change their positions only slightly. In a static case and a sufficient number of nodes FEM is a tool for precise calculations that can be very helpful for analyses of the final construction design. Actually, some companies use FEM to check the quality of the manually generated bores. This is especially necessary in die casting where blowholes – caused by insufficiently cooling – have to be avoided.

In manual bore layouts often only the mean temperature in the time period between injection and ejection is used to estimate the actually dynamic temperature distribution in the tool [Zö97]. Using only the static case simplifies the analysis a lot and helps to increase the calculations.

Due to the long evaluation times of the FEM, a fast surrogate model was introduced. The new model approximates the actual heat flow mechanism by modeling a heat radiation case. The exponential decrease of the cooling properties with increasing distance between bore and die surface is covered by both models. The quality of the FEM solutions can be reached in many cases. Only when a bore is shaded by an object or two bores are lying very near to each other, the approximation is not accurate. In most cases the shading effect and the superposition problem can be disregarded because of the strongly decreasing local character of the cooling effect. Furthermore, the superposition of the bores introduces an overrating of the cooling effect. This leads in the optimization process to bores with a larger distances to the die surface. The cooling distribution criterion introduces a separation of the bores and, hence, the overestimation error reduces exponentially again.

The basic idea of the radiation approach is to calculate the temperature distribution and intensity on the die surface radiated by a beaming cylinder. Each bore can be interpreted as a 'neon lamp' illuminating the die surface. Mathematically, each bore is modeled as a straight line between each two vertices of the bore circuit, P_i and P_{i+1} respectively, $i = 1, \dots, n$, having n bores. The radiation $\tau_{i,j}$ effecting each point M_j of the die surface can be calculated by integrating the radiation intensity of each point along the line between P_i and P_{i+1} . The radiation intensity of a light point decreases with $1/r^2$ with increasing distance r . Accordingly, the intensity of the line of points between P_i and P_{i+1} effecting M_j is the integral [WMM⁺04].

$$\tau_{i,j} = |P_i - P_{i+1}| \int_0^1 \frac{1}{(P_i + t \cdot (P_{i+1} - P_i) - M_j)^2} dt \quad (3.1)$$

The total cooling effect τ_j at point M_j generated by all the bores is the sum of all partial effects, i.e

$$\tau_j = \sum_i \tau_{i,j} \quad (3.2)$$

Each discrete point on the die surface M_j , $j = 1, \dots, m$ can be analyzed this way. Often a CAD surface is represented as a triangulation. In this case, the M_j can be taken as the centre point of each (small) triangle.

The global cooling effect f_l is calculated by the arithmetic average value of all τ_j . f_n is the normalized minimum effect value over all τ_j .

Function 3.2 can also be used to characterize the cooling distribution by calculating the statistical variances of the cooling values around each triangle τ_j . The global temperature distribution f_d is the mean of the variance values calculated over all triangles.

In the basic model a constant initial die surface temperature is assumed. In reality some areas are hotter than others, e.g. at the spray or where material accumulates. Demands for additional cooling is modeled by temperature values that are attached to the triangles. These cooling demands are integrated into the model (see Fig. 3.2).

The aggregated temperature effect f_t , which is used by the optimization process, is composed of three parts: the arithmetic average sum over all radiation contributions of all bores on all triangles (describing the absolute cooling effect) f_l , the temperature standard deviation (modeling the uniformity of the illumination) f_d , and the normalized minimum effect value f_n . All values are normalized and about similarly sensitive to changes in the bore design.




Figure 3.2: Cooling effect calculated via the radiation approach (left). Highlighted areas indicate better cooling than dark areas. Specific cooling demands (right) can be defined manually. Highlighted areas indicate higher demands than dark areas. Superposition (middle) of cooling effect and demand yields an effective cooling quality that is used for the optimization.

In multiobjective optimization it is possible to use a priori or a posteriori techniques. Both techniques have been used here. The a priori method is represented by an aggregation or scalarization technique. This method can be used adequately, if the relative weights of the contributing partial target functions can be defined a priori by the user. It is also known that linear aggregation should only be used in the convex case of the Pareto front. Actually, in our problem domain there are only very few practical cases known where the Pareto front is concave. Nevertheless, we also applied an a posteriori analysis of the Pareto front of the mold temperature control problem that proves this fact again, as shown later.

The aggregated fitness f is a function of the standardized restrictions f_{pen} , the manufacturing costs f_c (i.e. the length of the bores), and the temperature effect f_t :

$$f = f_t \cdot (1.0 + f_{pen}) \cdot (1.0 + d \cdot f_c) \quad (3.3)$$

$$f_{pen} = a \cdot f_{bc} + b \cdot f_{sc} + c \cdot f_{ot}$$

The all-over penalty f_{pen} is a linear combination of the number of collisions between two bores f_{bc} , the factor f_{sc} describes the collision between the bores and the tool surface and f_{ot} is a penalty for the bores that reach out of the tool. a , b , c , and d are arbitrary real value weighting factors.

The radiation approach has the advantage of very short calculation times. Compared to the FE method, which takes about 10 minutes for a simple geometry, the calculation time for analysis of the cooling design using the radiation technique takes just about 1 ms on a standard PC. Short evaluation times and good approximation results allow to efficiently utilize stochastic algorithms and to solve the complex problem of mold temperature control design.

Multiobjective Evolutionary Algorithms

For about five years the field of multiobjective optimization (see section 2.5.3) is in the focus of increasing interest. The classical single criterion optimization is a special subclass of the more general multiobjective optimization. In general, single criterion techniques cannot be applied to multiobjective cases without changes. Even the comparison of solutions is not as easy as in the scalar case, because sets of compromise solutions have to be compared with each other. Therefore, the optimization techniques have to be adapted to this class of problems. Optimization in this context means a gradual increase in quality for the solution sets, consisting of two or more partial target values. Furthermore, the calculated Pareto front, approximating the true Pareto set of the given problem, should be evenly spread at least in certain practically important regions. Additionally, of course, the approximation has to be calculated fast. In multiobjective optimization a lot of problems are very complex in an algorithmic sense, i.e. they cannot be solved in polynomial time. In combinatorial problems even the management of the solution sets can be NP-hard. Nonlinearity, a multimodal character of single or all functions, discontinuity and function ranges with no inclination, many nonlinear restrictions, high dimensions of the decision space and even dynamic features are characteristics of real world problems that make multiobjective optimization a very difficult field in mathematics [Meh05].

Most of the existing deterministic methods for multiobjective optimization are relatively old. Nearly all techniques find only one solution on the Pareto front per run. The theory in continuous multiobjective optimization is also not as well developed as in the single objective case, although there is very good literature available (see e.g. [Ehr00, Mie02]).

Multiobjective optimization techniques can be grouped into a priori, a posteriori and interactive techniques. Many approaches use the a priori approach, where all necessary knowledge about the fitness criteria and their relative importance is known in advance. In this case, typical approaches are scalarization of the multiobjective optimization problem via aggregation or lexicographical ordering. There are various aggregation methods that use e.g. weighting sums or sums of weighted factor combinations. Lexicographic max-ordering is an interesting alternative, because this technique allows to find points that are Pareto optimal [Ehr00]. This property is not fulfilled in linear weighting techniques. Today, a posteriori problems are generally solved by stochastic optimization as described later. Interactive methods are an interesting alternative to both extreme cases, where expert knowledge influences the search process directly [Meh05].

Deterministic techniques have favourable properties if the problem can be solved with these approaches. It is possible to give approximation estimations, and deterministic methods are also very fast. Of course it is not easy to find an efficient method for a specific problem. Furthermore, the NFL-theorem [WM97a] tells that there is no best optimization algorithm for any class of problems at all, although, reducing the general set of all problems to smaller classes, it is possible to find at least 'appetizers' [DJW98].

Looking for algorithms that can find good solutions and also cover large areas of the Pareto front needing only one run leads to multiobjective evolutionary algorithms. In the single criterion case evolutionary algorithms have proved their effectiveness in many cases, either in theoretical or in practical applications. They may not always find the theoretically best solution, but on the one hand in the multiobjective case it is difficult to prove the optimality of a solution anyway – especially if only simulations are available – and on the other hand in practice gaining a relevant improvement is often that difficult that already a certain melioration is highly desirable.

Multiobjective evolutionary algorithms (MOEA) are developing since the last decade [Gol89]. A relevant increase in research and application came with the introduction of the Pareto dominance based MOEA. Algorithms such as the NSGA-II [DAPM00b] or the SPEA 2 [ZLT01a] belong to the classics of the current state of the art. A description of current algorithms is available

in [CvVL02] or [Meh05]. A discussion of MOEA is also given in the chapter 'Challenges for Contemporary Evolutionary Algorithms' of this book.

Multiobjective evolutionary algorithms have the advantage that they do not have to be re-programmed when fitness functions are changed or new criteria or restrictions are introduced. They can be used either in scalar mode (classic single criterion EA) or in Pareto dominance mode (Pareto fronts and sets). Both approaches have been used for bore layout optimization.

In order to compare Pareto fronts with respect to their quality, various scalarization methods – so called metrics – have been introduced. For a well founded comparison see e.g. [vV99, CS03]. Metrics help to calculate a quality measure in the objective space of a Pareto front with respect to a given reference point, a known best Pareto front or a relative measure between two Pareto fronts, although the definition of the quality is not unique. Of course good approximations of the true Pareto front and well spread solutions are desired. Often the true Pareto front is not known. Only best solutions found so far may be at hand. Furthermore, practically obtained solution sets are discrete in nature. Metrics have to be able to cover all these multiobjective problems and to be as intuitive as possible. A technique that matches many of these issues is the attainment surfaces method [Kno02]. A variant of this method is the MMBBH-measure [MMBBH04]. This method has been used to assign a distribution measure and an approximation measure to one value which is needed for the comparison of algorithms. The solution sets for the comparison may come from different MOEAs, or the same MOEA with different parameter settings. This value was necessary for the application in the univariate design of experiments as used here to meta-optimize the MOEA parameter settings.

Statistical Design and Analysis

The model of a physical or technical system is generally only an approximation of reality. It has to cover certain properties and should help interpreting aspects of reality. Uncertainty, over-adaptation and generalization abilities are conflicting properties in the modeling and design of experiments that are necessary for the verification of theses or mathematical models. Some models can also be used for extrapolation in time, space or contexts, but most of them are restricted to a certain area of interest and use data sets as a basis to describe input-output relations. Usually only a restricted number of data is available, some data is missing or sparse, or the data contains errors or noise. In general, the design of adequate models is a very difficult task. This is especially true for the optimization in real-world applications because often no models, i.e. no mathematical functions, exist at all. The selection of the data for the assessment of a model is also non trivial in general, because its amount and distribution depends on the model and on the availability of the data.

A well established method for the analysis and optimization of problems, especially when the number of available data is low, is the statistical design of experiments. Design of experiments (DoE) methods need a minimum of data. The results can be used to derive functional dependencies between parameters and target values, and so support optimization purposes as well [WJ99]. Grid design plans have the disadvantage that the number of experiments grow strongly exponential with the number of dimensions of the parameter space. One-factor-at-a-time approaches are very popular but they do not allow to make statements about the interactions between parameters [WJ99]. Full factorial design plans use two or three parameters per dimension, fractional factorial design plans need even less experiments and allow to characterize factor interactions and optimization. The standard DoE models are linear or quadratic. Planar linear models, i.e. models without factor interactions, are used for screening. This means that in problems with high numbers of factors (parameters) the most relevant are filtered with respect to the influence on the responses (scalar outcome). In that phase Plackett-Burman design plans are used that can be calculated by systematic

schemes [Mon01]. In the modeling phase a linear model with interactions is used to describe interdependencies between factors, i.e. factors can influence each other multiplicatively. After the screening phase, DoE modeling only utilizes the relevant factors because the design plans are a little more expensive. Here factorial design plans are used that also follow schemes listed in the literature [Mon01]. Central composite design plans are typical for optimization [WJ99]. In the optimization phase of DoE, a quadratic model is used to fit the measured data via regression since this type of supposed dependency can be solved efficiently. Algorithms such as the conjugate gradients method are frequently successful.

Design and Analysis of Computer Experiments (DACE) [SWN03] is a technique for the generation of design plans and evaluation of experiments with deterministic behavior and quantitative character. The analysis of nominal factors, nondeterministic responses [BB05] and sequential introduction of new design points (SPO) [MMLBB05] are current state of the art extensions that improve this method.

DACE needs a space filling design. A typical plan is the latin hypercube design (LHD). Compared with DoE, LHD generates similarly few design points. The points are spread on a grid in a way that each column and line of the grid is occupied by one point only. Additionally, in the experiments the points are distributed in a maximum space filling fashion in the search space. The responses of the corresponding experiments following the LHD designs are interpolated by Kriging [Wik06a]. Kriging introduces a correlation between the measured points and yields a smooth interpolating response surface. The problem of noise in the responses can be solved by boot strapping [MMLBB05].

SPO (sequential parameter optimization) introduces an easy to use and efficient way to improve the design plans by iterative introduction of new design plans. New design points are set following a minimum trade off principle that minimizes the uncertainty of a prediction error (global search) and maximizes the local improvement to find better function values (local search) [SWJ98]. Starting from an optional DoE screening of the relevant factors, the first step in SPO is to perform usual DACE calculations. The Kriging model in the DACE approach plays the role of a statistical prediction tool. In the subsequent SPO iterations, only relatively few new points are introduced to improve the model fit step by step. Compared to the number of points used during the first SPO step, generally in the following steps only few (only 10% had been necessary in our experiments) new points are used.

The DACE/SPO technique is important when the number of fitness function evaluations has to be small. This is especially true when the evaluations are expensive. This technique can also be used for meta-optimization of EA parameters. Although the EA's behavior is relatively invariant against changes of the fitness functions, a tuning of the parameters can still help improving the quality of the results with a statistical significance.

Results

The evolutionary optimization of the model temperature control designs of casting dies starts with an initial definition of a set of possible solutions called a population. Each solution called individual is a vector of vertices of a polyline modelling the cooling circuit and the respective bore directions. The initial population contains values that are arbitrarily distributed in the volume of the tool. In the beginning it is possible that the bores may penetrate the surface of the die. These solutions have a relatively bad quality so these individuals are eliminated from the population during the evolutionary process rather soon. The initial parameters such as the number of bores, the geometric properties of the bores and the die, parameters that influence the optimization such as the population size, the selection pressure and many more values can be tuned with a graphical

user interface controlling the underlying long-running evolutionary computation process. Selected bore designs can be visualized online during the run and can be modified interactively. The system supports a re-evaluation of modified solutions in real time. This allows the user to watch the optimization process and to tune the parameters of the simulation according to the individual impression interactively. In the analysis standard die geometries were used. Alternatively the system can import, visualize and analyze STL-CAD-geometries. Due to the modular structure of the tool, the evaluation process can be linked to an evolution strategy or to other MOEA-tools. The interfaces between the optimization strategy, the visualization and the evaluation module are realized as an internet socket connection. Therefore, several system components can be exchanged with respect to the respective aims of research.

The standard ES has been modified slightly. Numerical experiments showed that using correlated mutations yield a slight improvement in the convergence properties. A significant melioration of convergence probability towards good solutions was realized by introducing a limited minimum step size that is declining slowly. This variation was useful because the standard step size adaptation proved too fast for this application. An increase in the freedom of search, i.e. using relatively large step sizes, helps to avoid getting stuck in local minima that are typical for the strongly restricted search space. Due to the fact that bores must not intersect each other, structures can get into constellations that cannot be resolved with small geometric step sizes. Large step sizes are an effective solution for this problem. Additionally the penalty functions for the intersection of fixed geometric restrictions such as pillars etc. that must not be intersected by bores can increase gradually in the simulation. This also improves the ability of the algorithm to explore the complete search space.

Uncertainties about the preferences of the user toward different quality properties can be modeled via a priori and a posteriori techniques. The ES approach aggregates the fitness functions using a product of individual target functions using normalized fitness values. This method allows to model fitness function interactions and helps to adapt the resulting scalar values according to the (subjectively motivated) preferences of the user. In this approach the single objective evolution strategy as described above has been used. This has the advantage that good initial parameter settings of the ES can be chosen according to good experience. The behavior of the ES is well known and the analysis of the optimization process can therefore be focused more on the convergence behavior of the bore structures toward good solutions. The fast fitness evaluation of the radiation approach is a necessary precondition for the efficient optimization using population based techniques. FEM analyses in the ES loop have been tested but the long evaluation times allowed only very small population sizes and few generations.

The Kit for Evolutionary Algorithms (KEA) contains current state of the art MOEA algorithms such as the NSGA II and SPEA 2, as well as two multiobjective particle swarm algorithms for treatment of multiobjective problems. KEA is written in Java and communicates via an internet socket interface with the evaluation tool. It provides a graphical user interface (GUI) for choosing between different MOEA algorithms and various fitness functions. The GUI also supports the tuning of the MOEA parameters. The Pareto fronts of two criteria multiobjective functions can be displayed online during an optimization run.

A typical optimization of the mold temperature control designs using the ES working on the standard test problem of a spherical die takes about five minutes on a standard PC. The MOEA optimizers are similarly fast. Of course in the MOEA approach not every Pareto-optimal solution (i.e. bore structure) is displayed online during the run. Single solutions have to be selected from the Pareto front (c.f. Fig. 3.3) by the user after the run and have to be displayed by the visualization tool of the evaluator separately. Figure 3.3 also shows that the Pareto front is constituted by clearly different domains defined by the number of bores. The general result is not unexpected: While




Figure 3.3: 200 solutions for runs with 3, 4, and 5 bores using the aggregating approach with an ES approximating a Pareto front. Inferior solutions relative to the Pareto front with respect to the two target functions are also displayed as they show the population's remaining variation span in the target functions space at a certain time of the respective evolutionary optimization run.

a larger number of bores will generally deliver better thermal performance it also causes more costs. A smaller number of bores will be relatively cost efficient but is strongly outperformed with respect to thermal efficiency by the other bore variants.

The MOEA optimizers are not limited to two criteria. The mold temperature control design optimization typically uses four quality criteria (cooling distribution, cooling intensity, costs) together with an aggregated value of the restrictions. Using a higher number of criteria than four or five – actually more than twelve criteria are implemented – is not advisable due to practical reasons. The Pareto front becomes much too complex to be analyzed by the human user even though there is no principal upper bound for an interactive Pareto front display as shown in section 3.3. In the literature on realistic real-world applications a maximum of seven criteria is suggested [CvVL02].

Looking at the test case's geometric results of the bore designs for the half sphere die surface, it is striking that a relatively high number of local optima appear. The solutions can be grouped to clusters of designs with a similar structure (see Fig. 3.4). The structures between the clusters differ significantly. This is an indication that the fitness landscape is multimodal. The average quality of the results of each cluster also differ. Even for a human expert it is not easy to tell what structure has a better quality than another. Therefore, the objective a posteriori evaluation is a valuable help for the designer to compare designs and to learn more about efficient structures.

The application of the standard statistical design of experiments proved to be an adequate technique to improve the strategy settings of the ES as well as of the MOEA, leading to superior target function qualities, i.e. better approximations to the true weighted absolute optimum or the Pareto front. The MOEA results have been compared using the MMBBH-measure [MMBBH04]. The meta-optimization using DoE assumes a quadratic behavior of the response values. In many practical applications and also near the optimum this assumption is very useful and an optimum can be found efficiently. Often an optimum lies on the border of the min-max-limits of the factors. Then the area of interest should be shifted to a more promising region and the experiments have to be repeated. Applying DoE to the ES and MOEA increased the quality of the results significantly, i.e. the values of the measure qualities improved more than ten percent compared to results with standard settings.




Figure 3.4: Three design solutions (first row) using each three, four, and five bores. Alternative solutions using the same number of bores (columns) and having similar cooling qualities are shown in the corresponding second row. Alternative solutions with the same (scalar) fitness values emphasize the fact that optimal designs of the cooling bores are generally not unique. Geometrically different solutions with similar cooling values are indicators for a multimodal fitness landscape.




Figure 3.5: DACE response fitness landscapes (grayscale) and errors (contour lines) for the parameter optimization of the ES. The best parameters are marked by a cross.

A more efficient approach to the design of experiments is provided by the DACE/SPO-method [MMLBB05]. Although also here extrapolation is not allowed, a larger region of interest can be scanned and modeled with free form surfaces (see Fig.3.5). The Kriging model shows very intuitively the interdependencies between the parameter settings and their influence on the quality. The strong stochastic character and the multimodal fitness function is a problem. DACE usually needs a deterministic behavior of the responses. The evolutionary algorithms generate a spectrum of results for the same parameter setting. The distribution of the experimental results can be handled efficiently with a bootstrapping method. The number of experiments per parameter setting can be kept surprisingly small, although the variance of the results is quite high. The

boot strapping method generates single representative response values that were used in the DACE model as deterministic responses. The number of additional new design points that are used to improve the DACE model in the SPO step is also quite small. Therefore, the DACE/SPO technique is very efficient and supports an intuitive interpretation due to the flexible Kriging visualization. Near quadratic dependencies between the parameters are displayed for the spherical test design problem. In DoE only quadratic hypotheses can be confirmed or rejected. The application of the DACE/SPO technique to the mold temperature control design optimization algorithm improved the EA results significantly.

Summary

The mold temperature design optimization problem shows a high structural complexity, the fitness functions are multimodal, complexly restricted, high dimensional and multiobjective and the uncertainty about the users preferences make the search for good solutions extremely difficult. An adequate geometric modeling as well as the objective modeling of qualitative and quantitative quality criteria is an important precondition for realistic real-world optimization. This problem is even more challenging when the evaluation of the criteria is time consuming and many evaluations have to be performed.

The modeling of the fitness functions follows an approximation of the real temperature flow evaluation using a radiation technique and the modeling of practical heuristics. The multiobjective optimization is performed via evolutionary algorithms. Different techniques such as scalarization and Pareto dominance were applied alternatively to provide different solution methods to the user. Speed, visualization and interactive optimization are important issues for real-applications. The systematic improvement and verification of the results have been done using statistical approaches such as classical design of experiments and modern DACE/SPO methods.

3.2 Decision support for intensity modulated radiotherapy planning

Karl-Heinz Küfer

Radiotherapy is, besides surgery, the most important treatment option in clinical oncology. It is used with both curative and palliative intention, either solely or in combination with surgery and chemotherapy. The vast majority of all radiotherapy patients is treated with high energetic photon beams. In this technique, the radiation is produced by a linear accelerator and delivered to the patient by several beams coming from different directions (see figure 3.6).




Figure 3.6: The irradiation gantry moves around the couch on which the patient lies. The couch position may also be changed to alter the beam directions.

In conventional conformal radiation therapy, only the outer shape of each beam can be smoothly adapted to the individual target volume. The intensity of the radiation throughout the beam's cross section is uniform or only modified by the use of pre-fabricated wedge filters. This, however, limits the possibilities to fit the shape of the resulting dose distribution in the tissue to the shape of the tumor, especially in the case of irregularly shaped non-convex targets like para-spinal tumors, prostate carcinoma located close to the rectum, or head-neck tumors in the proximity of the parotid glands and the spinal chord.

A significant advance in treating such difficult cases was the development of intensity modulated radiation therapy (IMRT) where the intensities on the cross-section of a beam can be varied. Using multi-leaf collimators (MLCs) (see figure 3.7), the intensity is modulated by uncovering parts of the beam only for individually chosen opening times and covering the rest of the beam opening by the collimator leaves. This permits shaping highly irregular dose distributions in the target volume.

An IMRT plan is physically characterized by the beam arrangement given by the angle of the couch relative to the gantry and the rotation angle of the gantry itself, and by the intensities on each beam. The treatment aim is - as a pragmatic definition - to deliver sufficient radiation to the tumor while sparing as much of the healthy tissue as possible. A major challenge in IMRT planning is to cope with these fuzzy goals and demands.




Figure 3.7: A Multileaf Collimator (MLC). The square opening in the center of the machine is partially covered by leaves, each of which can be individually moved. (picture from [SM01])

Concerning the beam arrangement, in most cases a concentric irradiation geometry is chosen: the beams meet in one iso-center as depicted in figure 3.8. Additionally, the beams are generally chosen to lie in the same plane, as positioning the couch to realize a full 3D geometry increases




Figure 3.8: Beam directions in a concentric irradiation geometry with an equiangular setup

treatment time substantially. Since the mathematical formulation of the planning problem including the irradiation geometry poses a much more difficult problem description¹, the beam directions are usually fixed manually. In most cases, there are fewer than 10 beams and often they are arranged in equiangular distance around the patient.

The remaining planning parameters - the description of the intensity modulations of the individual beams - may now be determined by mathematical optimization procedures. However, the choice of an appropriate objective function to optimize presents one of the biggest hurdles in formulating a mathematical optimization problem for IMRT planning. Many quality measures for a treatment plan based on different approaches have been proposed by medical physicists and oncologists. Common to all measures is that they are based on the distribution of the realized dose in the patient's body.

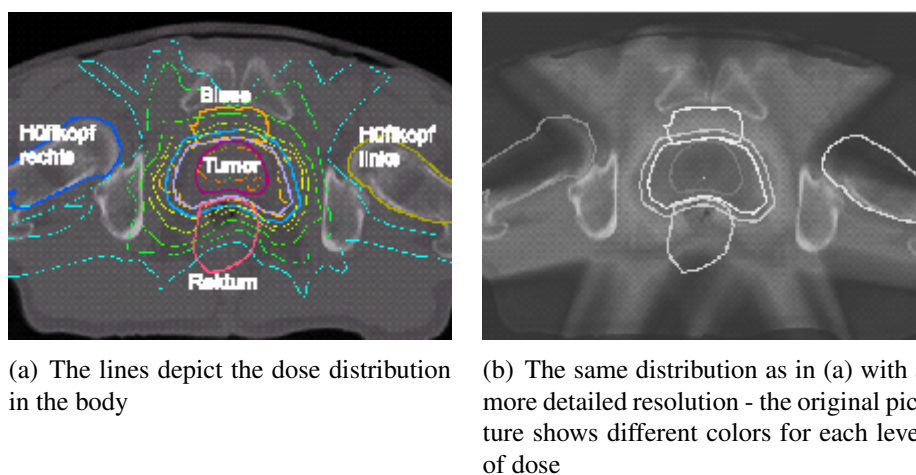


Figure 3.9: Visualization of the dose distribution in one of more than 100 CT-slices of a patient

Some measures typically used are, for example, the average dosages received by organs and tumor volumes. Often variations from ideal reference doses are measured. Yet another function is the variance around the mean dose in the target. Structurally, these commonly used quality measures are descriptive and comparable for the planner.

In most clinical cases, one or two objective functions are specified for the clinically relevant structures. This way, a multi-criteria formulation of the planning problem arises naturally. The quest for a solution that has a single optimal quality score is thus extended to a search for solutions that are *Pareto optimal* or *efficient*. An efficient treatment plan has the property that none of the individual criteria can be improved while at least maintaining the levels of the others. In our context, given a Pareto optimal treatment plan, lowering the dose in one organ can only be achieved if the dose in the tumor is also lowered, or at least one more organ receives a higher dose, for example. In general, there is a multitude of such solutions to any given multi-criteria optimization problem. We refer to all efficient solutions as the *Pareto set*.

The subjective nature of IMRT planning is in part due to the specific choice of objective functions by a treating oncologist. A different planner may prefer similar but different quality measures for a given case. Since all quality indicators in IMRT planning are designed to achieve the same fuzzy goal - allowing control over the spread of the tumor - they measure more or less similar effects. It may therefore safely be assumed that the quality measures are correlated with each other: good plans under one objective function are also good plans with respect to another

¹This problem is not convex and hence can not be guaranteed to be solved exactly in reasonable time.

objective function. In other words, a Pareto set obtained using one specific set of quality measures is most likely not altered too much if a different set of correlated quality indicators were used. This diminishes to some extent the severity of the fuzzy environment IMRT planning is situated in. As long as the planner uses quality indicators that correlate with commonly agreed on sensible measures, it is most likely that the treatment plans he obtains are of high clinical quality.

The software MIRA (Multi-criteria Interactive Radiotherapy planning Assistant) developed by the Fraunhofer institute for industrial mathematics (ITWM) in Kaiserslautern is designed to calculate Pareto sets and, more importantly, provides an interactive environment to select an efficient plan that is according to the notions of the oncologist. The latter is the implementation of a support system for a highly complex decision problem. Detecting planning possibilities and physical limitations in a database of efficient solutions while considering multiple clinical preferences at the same time is a very challenging task.

The optimization algorithms implemented in MIRA calculate a database of Pareto optimal treatment plans that cover a clinically meaningful range. As soon as the decision-maker has selected some quality measures, MIRA begins with the calculation of solutions. A high number of criteria leads to a relatively large Pareto set. Since this set is continuous, the possibility of displaying the infinitely many solutions is not an option. To still select a treatment from the Pareto set, it suffices to approximate the set appropriately. For this a so-called planning horizon is determined to exclude clinically irrelevant solutions. The Pareto set will contain a lot of plans where the trade-off between overdosage in the organs and a specific dose in the tumor will be too high to consider. MIRA deploys *extreme compromises* as corner points for the planning horizon in the Pareto set. They are the compromise of all possible combinations of the objective functions. The idea and its mathematical consequences are described in more detail in [KMS⁺05].

After this planning region has been marked, individual solutions are placed between the extreme compromises to approximate the Pareto set in more detail. Calculating an extreme compromise or another solution from the Pareto set requires solving a high dimensional optimization problem. Specialized optimization routines developed at the ITWM enable efficient calculations and creation of the databases in reasonable time. The strategy called the *adaptive clustering method* solves a sequence of approximate problems, which iteratively converge to the original problem. The solution of each previous problem is used to adaptively refine the approximation. The method is explained more fully in [SKM⁺04].

After the calculation of a database of treatment plans, the interaction with the decision-maker begins. MIRA has an interface called the *Navigator* (see figure 3.10) which enables the user to explore the Pareto set using visual controls. All user interaction with the Navigator are internally translated as optimization problems on the Pareto set and solved in real-time (see figure 3.11). This allows a smooth interaction with MIRA and explains the names “assistant” and “Navigator”.

The star on the left hand side is composed of axes for the different quality indicators. The objective functions associated with the organs at risk are combined into a radar plot, whereas the functions associated with tumor volumes are shown as separate axes. The interval on the axes corresponds to the range of values contained in the database for the individual objective functions. The white polygon marks the indicator function values for the currently selected plan. The shaded area represents the planning horizon.

The interaction with MIRA is characterized by two fundamental mechanisms which are patented for the use in IMRT planning by the ITWM:

- *restriction*, to alter the planning horizon, and
- *selection*, to alter the current solution.




Figure 3.10: The Navigator. The star on the left hand side shows the planning horizon and the current solution. On the right the dose visualization and other statistics for the current solution are shown.

Restriction allows the user to set the bounds of the planning region in both directions: to narrow it and to expand it. Unwanted parts of the Pareto set can be excluded from planning this way. The planning horizon displays a valuable piece of information for the decision-maker: the inner boundary close to the center of the star represents the highest ambition in an individual objective the user may have, while the outer boundary depicts the absolute highest “cost” in a criterion that has to be paid to improve another objective.

The line representing the currently selected plan has handles at each intersection with an axis and triangles for the tumor related axes. Both can be grabbed with the mouse and moved to carry out the selection mechanism. Internally, optimization routines ensure that the changes in the other objectives are at a minimum. This enables the greatest possible control over the navigation through the Pareto set: MIRA tries to change the current solution only in the direction the decision-maker wants it to.

The right hand side of the screen displays the corresponding plans concurrently. It is the visualization of the dose distribution parallel to the interaction with the Navigator that conveys the possibilities and limitations in a clinical case.

Several optimization problems are solved each second to allow a smooth interaction with the user. The sensitivity of the quality measures that are not actively changed to the objective that is currently altered is shown by the magnitude of the changes in the current solution while the handle is pulled. If small changes in one quality measure lead to rather large changes in some other criterion, the objective functions are particularly sensitive to each other in the region around the current solution. The decision-maker gathers a lot of insight and feeling for the clinical case using this valuable information and is thus able to accurately determine the trade-offs involved with each treatment plan.

We now demonstrate a small planning example for a head-neck case. These are typically difficult to plan, since the primary tumor can be located anywhere in the naso- and oropharyngeal area, and regularly the lymphatic nodal stations have to be irradiated because they are at risk of




Figure 3.11: A conceptual view on the functionality of the Navigator: while the user changes the current solution, the “distance” of the current solution on the surface to the point represented by a star is minimized. The line connecting the stars shows the path the user has dragged in the navigation by moving a slider. This is where the users would like to the solution to be. The shaded region depicts the planning region on which the path casts a “shadow”. The coordinates of this shadow are displayed in the Navigator.

containing microscopic tumor spread. This results in big, irregular shaped target volumes with several organs at risk nearby.

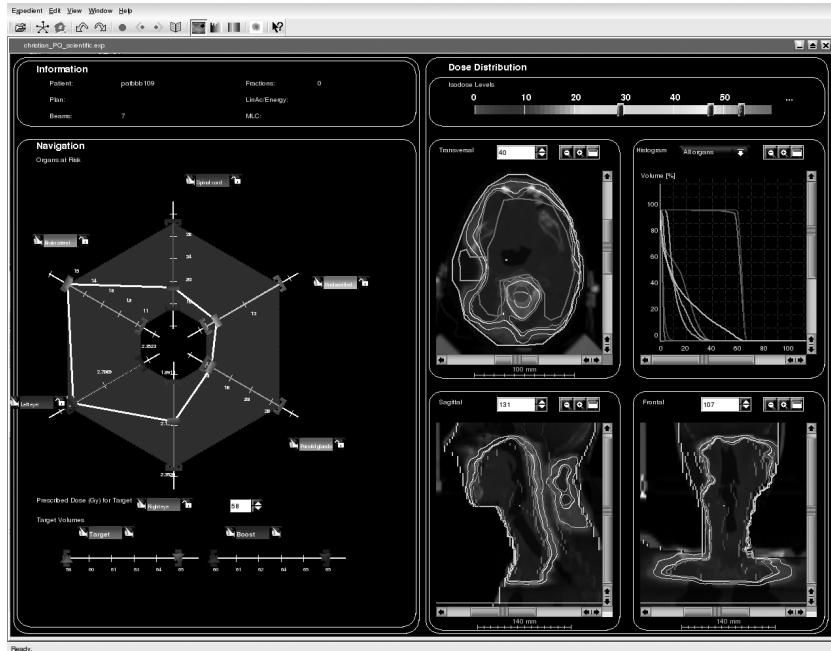
The salivary glands are such organs at risk that are quite radio-sensitive. The tolerance dose of the biggest salivary gland, the parotid gland, is approximately a mean dose of 26 Gy. The goal should be to spare at least one of the parotid glands. Otherwise the patient might suffer from xerostomia (a completely dry mouth) which can significantly reduce the quality of life. Other normal structures that have to be considered are (depending on the specific case) e.g. the brain stem, the spinal cord, the esophagus and the lungs.

The objective functions to two tumor volumes are the homogeneity of the dose around a specified mean. The axes of the star depict objective function values of the spinal chord, unclassified tissue, parotid glands, the left eye, the right eye and the brain stem. Now the interactive part of the planning process begins. The required compromises between the selected objectives can be evaluated by pulling the handles as described above. A first intuitive feel for the possibilities in the case is developed this way.

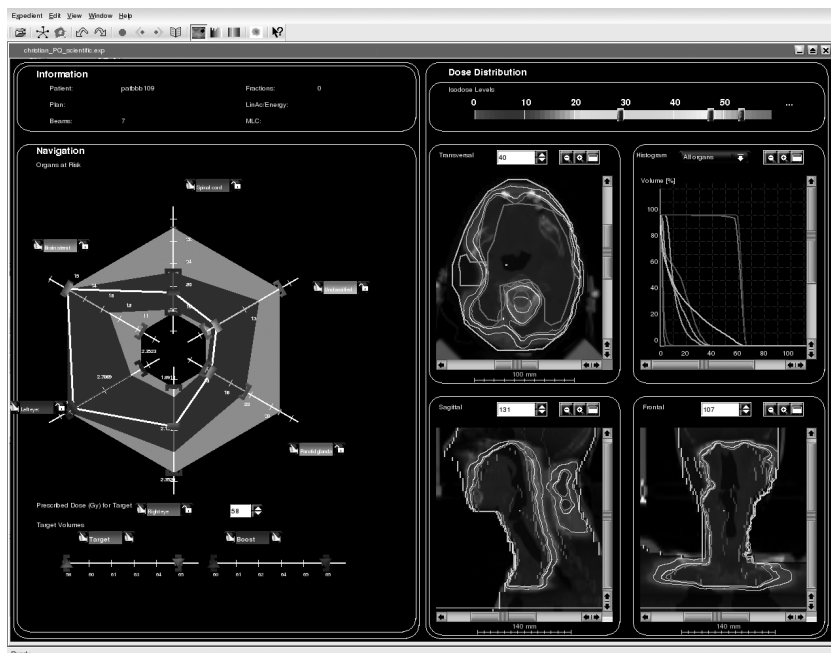
While the user alters the solution, the Pareto set and its dose distributions are explored in real time. MIRA also provides a locking mechanism for an organ. The user can fix a value on an axis, which will henceforth not be breached unless the lock is removed. All treatment plans that score worse than the fixed values are excluded from the planning region. This effect can be seen by the reduced planning horizon in 3.12(b).

Complex planning scenarios can be interactively explored in this manner and the best clinical treatment can be found within shortest time.

The design of the Navigator and the functionality of MIRA are targeted to spare the user



(a) Navigator at the beginning of the planning process



(b) The Navigator after some restrictions have been made - note the significant difference in the remaining planning domain.

Figure 3.12: The Navigator for the head and neck case

from mathematical formulations of optimization problems, and rather provide clinically descriptive information that is easy to interpret by experts. These aspects play a decisive role in acceptance of MIRA among treatment planners.

MIRA and the navigator present a tremendous opportunity for improvement of the planning workflow in hospitals all over the world. The graphical user interface for navigating a complex solution set and the flexibility in modeling provided by MIRA also facilitate interaction between physicians and medical physicists during the planning stage.

All optimization processes in commercially available planning packages for IMRT are based on mixing the objective functions if they even provide the opportunity for adding more than one. No currently available software provides multi-criteria decision support to the extent that MIRA can realize.

3.3 Multi-criteriality in the public perception: Experiences with a multicriterial decision support by an interactive program

P. Roosen and K.-H. Küfer

How do 'uneducated, ordinary people' tackle multi-criterial decision making? What strategies do they pursue to 'obtain what they want'? Professionals regularly and explicitly dealing with these kind of problems know several methods of approach, each of those being thoroughly investigated and assessed with respect of their predications and limitations (cf. section 2.5.3). In spite of the frequency of such problems also in everyday life this knowledge is not widely spread, though, so it is interesting to investigate whether non-professionals can make use of a support tool oriented towards this decision domain. The general idea is that the last resort of objectivity is the Pareto set of a given problems, so its display will yield the best possible decision support with respect to a self-restrained potential influencing of the decider by the tool provider.

While most practical problems present more than two target qualities to be optimized there is no static presentation tool that would provide a intelligible, ideally printable, selection display for the decider to work on. Therefore a computer-based tool was created to support interactive, visual pareto set testing poking. As the principal display mode is based on a star plot, the system was called 'Pareto Star' accordingly. Its applicability is not limited to pure pareto sets, though, as in practical applications there is not necessarily a sharp distinction of important decision criteria vs. irrelevant ones. Instead, a kind of continuous transition of criteria more or less crucial for the decision is frequently observed.

The designed decision support system is not a tool to *identify* a pareto set — it requires the existence of a database of realizable solutions with their individual target function values. There are cases, though, that exhibit continuous settings or operation modes and hence a dense pareto set of unlimited member count, so no distinct, objectively chosen points in target function space can be provided. Here the principle of uncertainty should be kept in mind, however: Small changes in the relative trade-off target settings usually do not make large differences in the eyes of the decider. So it suffices to identify sensible representants of the originally dense Pareto set and present them as individual target function (and accordingly configuration parameter) combinations in the selection tool.

As a practical test bed the selection of an handheld GPS receiver was chosen. There are several practical selection criteria, and no all-purpose devices serving the whole range of private navigation tasks exist. So the potential buyer is confronted with the selection problem, being an ubiquitous obstacle when scanning respective electronic discussion fora for answers and opinions. The tool devised as decision support is implemented as a web service, extendable to almost arbitrary multi-criterial decision problems. It operates on three basic tables, implemented by means of a database system: i) the list of potentially interesting criteria, ii) the list of available devices, and iii) the cross-combination table of their individual characteristics.

For the given decision domain there are about 20 parameters potentially relevant to different usages. Those may be as different as navigating a sensible automotive way in the streets, with upcoming traffic jam information being considered in a dynamic re-routing along the way, and trying to find small treasury caches somewhere in nature as recreative outdoor activity [Gro05]. As the respective items tend to be relatively costly (ranges in autumn 2005 between about 200 and 1000 EUR) many users try to cover different application domains by choosing just one device, sensibly supporting potentially conflicting usage scenarios. This in turn individualizes the target function weighting to an always new set.

Decision support tool setup

The decision support system was set up as a (German) web service, to reach widely distributed people as best as possible. It was placed on a non-commercial, topic related website of one of the authors with sufficient traffic to be found by chance. Additionally it was announced in various fora concerned with recreational GPS use, to attract the largest number of people possible. As there is no chance of forwarding a detailed usage manual to the casual visitor, the entry page was set up as such (fig. 3.13).




Figure 3.13: Introductory page for the 'Pareto Star' GPS handheld receiver selection aid. A small explanation of the usage is given.

In a very brief explanation the user is informed how to choose his favorite target values and how to maneuver in the star plot. Finally a hint is given that the resulting list of complying receivers is available via the central number in the star, indicating how much of the basically available devices are still left.

Starting the application by clicking the appropriate link ("Los geht's") leads to the personal criteria definition page (fig. 3.14). Here the user is requested to select at least three criteria that he regards as important. For every criterion a short information box, shortly explaining what is evaluated and which possible values can appear for it, pops up if the user moves his mouse pointer on its name. After pressing the 'Aktualisieren' (update) button the first star plot is presented (fig. 3.15).

The axes of the plot are directed so that the better values always lie in the center of the star. (Non-orderable properties, like housing colors or geometries, cannot be handled within the scope of the tool presently.) The different colors of the plot areas indicate different things. On any axis there are transitions from light blue to yellow to red to white. The transition from light blue to yellow indicates an individually set limit of acceptance (see below) that the user tolerates for that criterion. Without any prior interaction that limit is automatically set to the lowest value the respective axis can display for the given set of choices. The consecuting yellow part on any axis

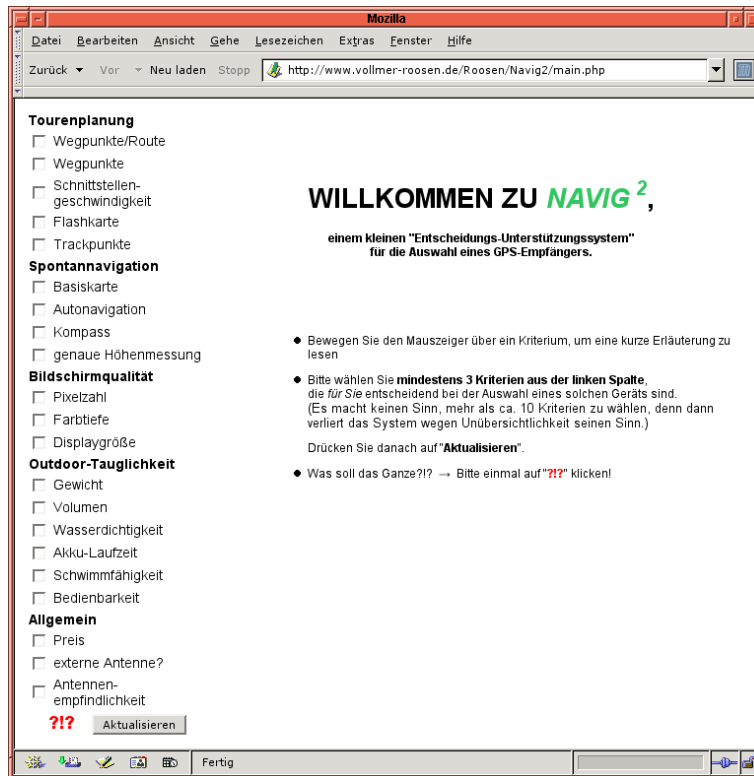


Figure 3.14: When entering the GPS receiver selection aid the user is asked for his personally most important properties of the devices. A small step-by-step guide is presented to help him starting. By clicking on the blinking '?!?' he obtains information on the background of the service.

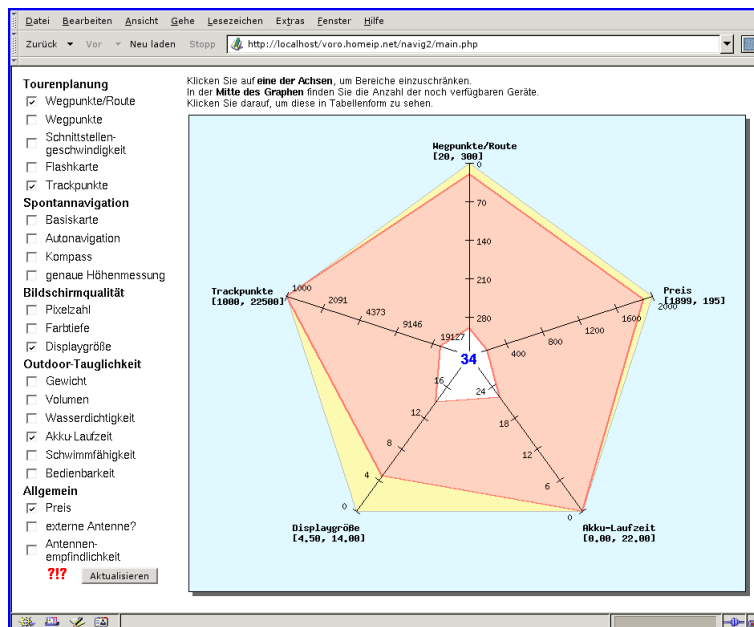


Figure 3.15: After entering at least three personally relevant criteria. In this example the five criteria 'waypoints per route', 'track points', 'display size', 'accumulator runtime' and 'price' were chosen.

indicates the range of values that is not covered by the choices' list, without being limited by a user's action though.

The red part on the axis is the range that the real properties of the discussed GPS devices cover.

As the display represents *all* devices not excluded by an individual acceptance limit, only the range for the whole set, not the individual properties of the devices, are displayed. The premise here is that the decider should concentrate on his basic individual demands, not on the concrete criteria of the devices.

The white range on any axis finally indicates a range that is not or no more accessible, either due to generally limited performance in the whole set of choices, or to a limitation that was defined for other criteria by the decider.

In the center of the star plot there is a thick blue number. It displays the number of choices left when the Pareto Star system honors the limitations set by the decider. For the GPS handheld devices there is an unlimited number of 34. The definition of a personal acceptance limit for any criterion selected beforehand will reduce the count of remaining devices respectively. The limits of the other target functions will change as well, both for the lower and the upper bound of the red range, as is evident when investigating the principally counteraction physical properties of (largest possible) display size and (smallest possible) device weight. Specifying a low value for the maximally accepted weight will reduce the maximum available screen sizes of the remaining items.

The point and click procedure on the accepted values, representing immediately the preferences of the the decision taker, is re-iterated for arbitrary axes until the number of remaining items is reduced to a tolerable quantity, quite often around 5 to 8. During this process the decision seeker may as well drop criteria from his selection as choose additional ones from the continuously displayed left criteria list column. Choosing an additional one will initialize another axis without a set boundary value, but obeying limitations on its criterion range imposed by set limitations on other criteria. After that the central number should be clicked to obtain a simple table of those devices, listing *all* their criteria properties (fig. 3.16).

Bild						
	Großes Bild	Großes Bild	Großes Bild	Großes Bild	Großes Bild	Großes Bild
Name	Etrex Legend	GPS 72	GPSmap 60C	GPSmap 60CS	Meridian Color	Meridian Gold
Hersteller	Garmin	Garmin	Garmin	Garmin	Magellan	Magellan
Wegpunkte/Route	125	50	250	250	50	50
Wegpunkte	1000	500	1000	1000	500	500
Schnittstellen-geschwindigkeit	115200	115200	1e+06	1e+06	115200	115200
Flashkarte	0	0	0	0	4	4
Trackpunkte	10000	2047	10000	10000	2000	2000
Gewicht	150	215	200	200	240	240
Volumen	171	390	312	312	378	378
Wasserdichtigkeit	7	7	7	7	7	7
Akku-Laufzeit	18	16	20	20	14	14
Schwimmfähigkeit	0	1	0	0	1	1
Bedienbarkeit	1	2	2	2	2	2
Preis	289	259	559	639	578	438
externe Antenne?	0	0	1	1	1	0
Antennen-empfindlichkeit	0	1	1	1	1	1
Dominierende Geräte				GPSmap 60C	Meridian Gold, Meridian Platinum	

Figure 3.16: Display of the properties' list of the remaining items, after reducing their number via the interaction with the star diagram. A combined image is shown because the computer screen is not large enough to contain the table as a whole.

Besides the potential display of the criteria description popups there are additional functions and bits of information in this list. First of all it may be sorted in ascending or descending order for any criterion by clicking on the respective small arrows on the left of the criterion's name. This supports the notion that usually there is a dominant criterion (mostly: the price) even if a multitude of relevant ones has been selected before.

Secondly the columns of the list are colored differently. Green columns indicate pareto-optimal items *with respect to the criteria set and the individual accepted ranges settings of the decision seeker*. It must be stressed that this pareto set is highly individual, as a change in the criteria list or alternative thresholds of acceptance may significantly change it. Columns colored white & blue indicate that the respective item is not a pareto set member but still complies with the imposed limits. In each non-pareto column the dominating devices are shown at the bottom of the table (row 'dominierende Geräte'), so the decider can immediately compare them with each other. The reason to include the non-pareto optimal devices in the list is the lack of a practical sharp distinction between relevant and irrelevant criteria in the selection process. Even if a criterion has not been included into the interactive Pareto Star selection that does not indicate that it's value is completely irrelevant. If, for a given device, such a factor becomes very favorable while the chosen 'important' ones strongly resemble each other it may as well serve as a secondary, hierarchically less stringent decision argument. This leads to a discussion on parity and hierarchy of optimization goals and will be revived in a later section.

Usage analysis

The decision support system was designed in order to monitor 'uneducated' users' reactions on the availability of an individually determined pareto set indicating filtering system since this kind of decision domain is ubiquitous in normal life. To attract a sufficient number of users the tool was placed on a well-frequented web site dealing with the application of smaller GPS devices. By means of small introductory notes it was announced on several related web-based fora.

On the web site itself a visually quite dominant teaser (red, bold and blinking '?!?' click field, placed directly next to the main display refresh button) was placed to motivate users to express direct responses. If a user clicks it he becomes a short note on the background of the tool and a text entry field to submit his impressions on the tool. In addition to those (rather few) direct responses the usage of the system was monitored by analyzing the server log files over an interval of about 3.5 months.

About 30 direct responses were received. Half of them just helped correct some faulty decision data on the ever-so-changing specifications of the individual receivers (due to frequent firmware revisions that the manufacturers provide for their customers). Those responses indicated that the users observed the provided data with thorough scrutiny and were eager to clean up the data base in the sense of a community effort. The other half stated that they thought the tool really helped deciding on a device to be bought. Some of those responders remarked, though, that they already bought their devices prior to finding the decision support tool. The retrospection of their purchase decision was mixed: Some stated they might have decided otherwise if they had known about the Pareto Star beforehand, some reported an affirmation of their stomach based former decision. No response discussed the tool and the method by itself, though. Obviously the content aspect was too dominant to focus on the methodology.

In order to obtain more of those methodological data the server log file have been analyzed. Due to the setup of the service as LAMP system (containing the components Linux, Apache, PHP, MySQL) and the input data submission mostly via the http GET method most of the user settings

with respect to choice and number of criteria could be reconstructed (table 3.1). The total number of 2600 visitors should provide an adequate sample size to consider the data as trustworthy.

Number of visitors (total)	2600
→ number per day	24
Star interactions per visit	2.7
Result table requests per visit	1.8
Mean number of criteria chosen for the star diagram	5.5

Table 3.1: Statistical access values for the Pareto Star decision support system for GPS handheld receivers.

The anticipated essence of the tool — the toying with the star diagram — was used rather sparsely with the averaged 2.7 interactions per visit. Here a definitely larger number was anticipated. There is an explanation, however, in the supposition that the users tend to be somewhat result-oriented (in the sense that they want to decide on a concrete device) and that the playing around with the tool did not satisfy their immediate curiosity just *which* devices were still in the race and which had already been left behind.

Another point may be the non-familiarity of the users with the density distribution of the individual criteria. If a scale indicates that a large range of values is available it does not show that perhaps 95 % of the devices are clumped in a (possibly non-favorable) small part of the span. This soon leads to rather sparse result lists. The average number of less than two result table requests per visit once again shows that the users do not tend to play with the tool, but instead are very target-oriented. After receiving a result table they rather seldom reiterate the star plot selection process with changed criteria settings.

This in turn may be interpreted as result of economizing the selection process in the sense of a frugal decision method: The users seem to obtain a convincing final selection list by just focusing on the first result table. The fact that usually familiar brands and types are listed will enhance that notion, once again supported by the Bounded Rationality view of typical human selection methods, here 'take the best known'.

These interpretations, derived from the average system's usage, do not hold for some very eager users who indeed played around quite intensely with the tool, reaching well above ten or fifteen interactions with the star diagram and a respective number of result lists.

Real-world applications

The developers presented the tool to several hardware and service providers in very different application domains, like used automobiles reselling, real estate brokerage, recreational accomodation services, or last minute flight ticketing. In almost every discussion the tool was honored as a very interesting method of information retrieval and assessment. Nevertheless the tool was not used for a respective innovative information preparation because a conflict of interests was regularly perceived. This may well be explained by discussing a potential application in the automotive reselling business.

For an internet-based vehicle reseller the usual business model is to provide a presentation platform for an arbitrary number of private vendors that publish their offers on it. For larger sites well above a million vehicles are offered at a given time. So even if the list size of criteria usually deemed relevant is estimated to five or six there are lots of offers not nearly approaching pareto optimality in the sense of the potential purchasers. If the trade platform offers a method to easily identify the most interesting ones in the sense of constructing a pareto front, most offers would

never (or at least only very rarely) be shown. Vendors intending to use the platform would turn away from using it. This in turn conflicts the business model of the web site provider who would be the one to offer the additional pareto set filtering service — hence the aversion to put such a tool to work. It may be guessed that most potential vendors have a certain kind of qualitative notion about the competitiveness of their offers, so the availability of a selection tool would shy them away to other platforms *not* offering it.

This situation holds for most points of electronic sales, therefore it is not to be expected that this methodology of multi-criterial decision support will find its penetration into a broader usage. More promising fields of application are areas where all participating players do have a sincere interest in an objective target criteria assessment. One of them is presumably the consumer consulting service that does not earn its money by selling items but with offering the best possible objectivity in consumer information requests.

3.3.1 Refinement of the selection tool

The experiences with the usage behaviour led to the conclusion to redesign the tool. It should both familiarize people with a multi-criterial selection process support and better meet their expectations with respect to their desire to obtain almost immediate practical suggestions. The new 'Pareto Star' system is consequently not a star plot diagram any more, but rather a semi-hierarchical table of bar charts with similar, but enhanced colorization ranges relative to the existing implementation.

While the final result table display will not be changed very much, the multi-criterial selection process is significantly changed towards a faster 'real results' display, partially implying a partially hierarchical expectation concept. Starting with an empty list of selection-relevant criteria the user will be asked to choose a first criterion (fig. 3.17).



Figure 3.17: New multi-criterial decision support system with instantaneous semi-hierarchical result display. (The visual appearance of the interface itself represents an intermediate development phase.)

After selection of one or several personally relevant criteria the result bars (fig. 3.17, left side) display a differentiated colourization (see below). A set of adequate receivers is presented on the right side of the interaction window, being background-colored in light orange to display a Pareto set membership as in the older version of the Pareto star system final result table.

The order of parameter bars on the left side is treated hierarchically, with the importance of the criteria decreasing from top to bottom. Accordingly the order of actual receiver displays is sorted in a strictly hierarchical way. Receivers exhibiting identical settings for the most important goal are sorted by the second important one and so on. In case of an additionally chosen criterion the list on the right is resorted accordingly. It is as well possible to switch places and hence the order of importance for the selected target criteria, with a respective effect on the order of the actual receiver display.

The graphical bars describing the single criteria contain additional information relative to the Pareto Star. Favourable settings are always on the right side of the bars, so the user will tend to click into the bar to limit the respective value with increasing rigidity from left to right. His personal limit is indicated by a thick black vertical bar in the horizontal criterion distribution, with the chosen value being displayed on top of it. On the left of that marker there is a white background. On the right usually a light green range is situated, indicating an interval that does not contain items with respective values. A light blue range follows, hinting towards actually available items. Contrary to the Pareto Star an additional background colored range follows: The red section indicates the interval of target value settings that will lead to a complete emptiness of the receiver list. This range is influenced by the user's limit settings in other criteria, due to the fact that those restrict the available choice in others.

In the former Pareto star system the user did not receive information about the density and distribution of the criteria values independent of his personal choices. Therefore an additional piece of information was introduced in the new hierarchical system: Small black dots in the center of the horizontal bar indicate the actually available values in the underlying database.

In the end the user will choose to display the resulting items table that is presented in a similar fashion as in the presently available Pareto Star system.

Finally, the new system contains visual procedures to work on discrete and unsortable list criteria. Usually there are values that cannot be ordered by their values, like the color of a car body, or the selection of edge shapes for a chair or a desk. Here the user will be able to define his or her personal preferences, with the displayed order of potential values reordered upon each new change of choice, separating the desired ones from the undesired. As soon as this selection system is in a stable condition it will replace the present Pareto Star, and its availability will be announced once again in the respective fora to attract new users.

3.4 Molecular Biology: RNA Structure Optimization on Large Neutral Networks

Peter Schuster

Structure formation in biology is basic to life and at the same time the prototype for molecular complexity. Folding biopolymers into stable structures is visualized as an optimization process that commonly follows the criterion of minimizing free energy. Function in nature is designed and optimized by evolution, which can be understood – in form of the Darwinian mechanism – as an optimization of a cost function. In natural ecosystems, in particular when they are free of human intervention, the cost function for evolution is called fitness and measures the mean number of fertile progeny in the next generation. The Darwinian mechanism operates in populations with a dispersion of fitness values: Each variant with a fitness larger than average will be present more frequently in future generations whereas the frequencies of all variants with fitness values below average will decrease. Fitness counts offspring only and does not deal with the complexity of the objects, which are optimized. Hence, the rules for Darwinian evolution are the same whether we are dealing with molecules, with organisms, with groups of individuals or with any other variable object.

Optimization of molecular properties by means of evolutionary methods based on variation and selection (SELEX)² is an established technique in biotechnology. Out of all classes of biomolecules ribonucleic acids (RNA) are best suited for the evolutionary design of products with predefined properties, because they can be easily amplified either directly by replication or via reverse transcription (RNA→DNA), DNA amplification, and transcription (DNA→RNA). The latter three step process is often applied because DNA amplification is easily achieved by means of the polymerase chain reaction (PCR). In addition, RNA molecules have a richer repertoire of functions than their DNA analogues. Various classes of RNA molecules with specific properties have been produced by evolutionary techniques. We mention here three examples: (i) Molecules that bind specifically and with high affinity to targets called *aptamers* [Klu06], (ii) molecules that catalyze biochemical and chemical reactions called *ribozymes* [WS99], and (iii) molecules that have more than one long lived structure called *riboswitches* [WB03, VRMG04]. Aptamers in particular those with high specificities and affinities are very useful tools for highly specific detection of the presence of their target molecules and are used for analytical and diagnostic purposes (For an example of an aptamer discriminating between the two closely related molecules caffeine and theophylline see e.g. [JGPP94]). Ribozymes were discovered first in nature and later on produced in the laboratory through modification of RNA molecules by variation and selection. Alternatively they were also evolved from molecules with random sequences. The repertoire of chemical reactions, which are known to be catalyzed by ribozymes is incredibly rich. Riboswitches, finally, are RNA molecules with two or more long-lived (meta)stable³ conformations. RNA switches may be self-regulated or the conformational change may be triggered by a small molecule binding to the RNA. In nature the small molecules are often metabolites, which regulate the pathway leading to them through inactivating translation via a conformational change in the messenger-RNA.

²SELEX stands for selection by exponential amplification [ES90, TG90].

³Metastable states are local minima of the conformational energy surface. They have a characteristic life time that is determined by the height of the barrier connecting the state with a conformation of lower free energy. RNA molecules having a long-lived metastable state in addition to the minimum free energy conformation are commonly called self-regulated RNA switches.

3.4.1 The Darwinian mechanism

Optimization in the sense of Charles Darwin as laid down in his seminal book on the origin of species [Dar59] is a universal phenomenon taking place in populations. It can be casted into the interplay of three processes:

- (i) multiplication, in particular reproduction of organisms or replication of molecules,
- (ii) variation in the form of mutation or recombination, and
- (iii) selection as a consequence of limited resources providing upper bounds to population sizes.

Since the conditions sustaining these three processes can be fulfilled likewise by polynucleotide molecules in cell-free assays, cells or multicellular organisms, evolution in the sense of Darwin occurs in populations of very simple as well as well as very complex objects. Multiplication, variation, and selection take place in specific environments to which the populations adapt through a Darwinian process. Studying evolution is commonly complicated by two problems: (i) The times required for selection and adaptation in populations of higher organisms that have generation times from weeks to years are too long for direct observations, and (ii) changing environments make the situation very complex. This is particularly true for evolution in nature where sometimes large numbers of species evolve together or coevolve in the same ecosystem.

Encoded information on the reproduction of individuals is stored in the genotype⁴ being a DNA or RNA molecule. The genotype is unfolded to yield the phenotype which is imagined best as the fully developed individual in its appearance and with all its functions and properties. Success and efficiency of evolutionary optimization are based on different roles played by genotype and phenotype. Variation operates on the genome (figure 3.19) whereas selection operates on phenotypes exclusively. What counts in the selection process is the number of (viable and fertile) offspring in the next generation. This quantity is the fitness, f . In nature, and particularly for higher organisms, fitness is a highly complex function of the phenotype and the environment, which also includes the other organisms living in the same ecosystem. All species evolve together, co-evolution superimposes its dynamics upon selection and populations may never reach their optimal states. Evolution experiments were conceived for constant environments that reduce this complexity. In test-tube evolution of RNA molecules further simplification is possible: Genotype and phenotype are different properties of the same molecular species – sequence and structure, respectively – and fitness can be reduced to an (over-all) rate parameter under suitable conditions.

The dichotomy of genotype and phenotype is fundamental for the success of Darwinian optimization. It introduces a random element into the evolutionary process since variation of the genotype and evaluation of the phenotype are uncorrelated.⁵ In other words a mutation does not occur more frequently because its phenotype has a selective advantage. Similarly the random element is indispensable for the Metropolis search method since it enables the algorithm to scan the entire solution space.

3.4.2 Evolution in the laboratory

In order to render the evolutionary process easier intelligible optimization has been studied under simplified and controlled conditions. Three examples are mentioned here: (i) Bacterial evolution

⁴In molecular genetics the genotype is called the *genome*.

⁵Mutations are often considered as random. This is not entirely correct since an individual mutation step is a chemical reaction as correct replication is. Small rate parameters are common for mutations and then the occurrence of a particular mutation falls into the domain of stochastic events.




Figure 3.18: **Basic mechanisms of RNA replication.** RNA replication commonly follows a complementary mechanism: A double-helical (plus-minus) duplex is synthesized from a (single) plus strand by making use of the complementarity of Watson-Crick base pairs. The critical step in replication is the separation of the duplex into two single strands because long double helical stretches are bound strongly. Then dissociation of double strands requires raising of temperature. In RNA evolution experiments separation into single strand is fulfilled by the replicase that prevents the formation of long double helical stretches by separating them into the two single strands that form their own structures.

[ECL96, PSM⁺99, EL03], (ii) an example of arms race in virus evolution [WWB05] and (iii) evolution of RNA molecules in serial transfer and automated evolution machine experiments [Spi71, MPS67, BG97, SE97]. The principle known as serial transfer is the same in both types of experiments: Growth medium⁶ is prepared and infected, either with bacteria or with suitable RNA molecules. Then, reproduction starts instantaneously and after a defined time interval (Δt) during which the population has grown to a certain size a small fraction of the sample is transferred into fresh growth medium. The procedure is repeated and continued until evolutionary phenomena can be detected. In the experiment with bacteria a transfer was made every day and the series runs now over more than twenty years already. In RNA test tube evolution the number of transfers is typically around one hundred.

Bacterial evolution is dealing with whole cells whose genotype is a DNA molecule with a few million digits called nucleotides taken from a four letter alphabet, $\mathcal{A} = \{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{T}\}$, and attached to a one-dimensional periodic 2'-deoxyribose-phosphate backbone. The phenotype is the bacterial cell with its highly complex structures and functions. Two results obtained in the experiments carried out with *Escherichia coli* bacteria in the laboratory of Richard Lenski are important for the comparison with the evolution of molecules: (i) Changes in the phenotype occur stepwise

⁶The medium is either a nutrient for bacteria, agar-agar for example, or a solution containing everything that is needed for replication.




Figure 3.19: **Basic mechanisms of sequence variation.** The upper part of the figure sketches three classes of mutations: (i) A point mutations where a single digit is changed and the sequence length remains constant, (ii) an insertion where part of the sequence is replicated twice, and (iii) a deletion where part of the sequence is not replicated. In the lower part we show a case of symmetric recombination between two sequences of equal lengths leading to two new sequences both with the same number of nucleotides. Asymmetric recombination (not shown) leads to sequences of different chain lengths.

and not gradual [ECL96]. The phenotypic changes concern, for example, the size of the bacteria, which increases in steps under the laboratory conditions. The larger variants reproduce faster and grow out the smaller ones. (ii) Changes in the DNA sequence caused by point mutations being single digit exchanges progress continuously in the sense that the rate of change is a constant [PSM⁺99]. A strong indication for adaptation to the conditions of the serial transfer was derived recently: Parallel experiments were carried out in twelve populations or clones that were recorded over 20 000 generations [CRL03]. Comparison with the ancestors showed characteristic changes in gene expression data⁷ but roughly the same increase in fitness. Eight out of the twelve parallel experiments revealed mutations in one key regulatory gene. On the genomic level, however, no two point mutations were identical.

Out of many experiments that have been performed on viruses one recent study performed in the laboratory of James Bull is remarkable [WWB05], because it provides a laboratory case study on co-evolution of attack and defense mechanisms in the sense of arms race. A population of DNA bacteriophages ϕ X174 has been grown for 180 days in a suspension of *Escherichia coli* bacteria corresponding to 13 000 phage generations. This phage has a small genome of 5 386 nucleotides, which is only one thousandth part of the genome of the bacterial host. Nevertheless, the phage parasite is able to cope with the evolving defense mechanism of the bacterial host. This experiment is an excellent example of co-evolutions under laboratory conditions. The authors identified mutations in the phage genome that increase its virulence. The arms race will continue until one of the two partners reaches the end of its capacity to adapt.

⁷In cells not all genes are active or *expressed* that means transcribed and translated into protein. Mutations in regulatory genes may lead to different activities of the corresponding regulatory proteins, which in turn modify the availability of structural or metabolic genes.

Evolution of RNA molecules in the test tube was studied first by Sol Spiegelman and his group [MPS67]. The nucleotide sequences of RNA molecules are the genotypes. In particular, they are strings built from the four letter alphabet, $\mathcal{A} = \{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{U}\}$,⁸ attached to a backbone that is almost the same as in the case of DNA. The lengths of RNA molecules with functions in nature genotypes vary from between twenty to about thirty thousand. Small interfering RNAs with regulatory functions are commonly twenty to thirty nucleotides long, artificially selected small aptamers have about the same lengths, and RNA molecules with catalytic functions derived through selection experiments have chain lengths from one hundred to several hundreds [WS99, Jäs01, Klu06]. RNA genomes in nature range from several hundred in viroids [FHM⁺06] up to thirty thousand in the largest known RNA viruses [GEZS06].

The phenotype in evolution experiments with RNA molecules is the spatial structure of the molecule, which is the carrier of all its functions and fitness relevant properties. In Spiegelman's experiments selection is optimizing replication rates, which are indeed the molecular equivalents to fitness values. In this way these experiments correspond to natural selection where fecundity and viability of offspring, subsumed together in the fitness values, are the only criteria for survival. Darwin himself used the results of animal breeding and plant crossings as an important support of his theory of evolution. In this case survival is not only a result of unguided fitness because human intervention introduces new criteria into selection. The molecular counterpart of animal breeding is directed evolution. It is carried out with the goal to optimize other properties than replication rates. Such properties are, for example, binding affinities, catalytic efficiencies, specificities, thermodynamic stability or robustness against mutation. Optimization following such criteria requires a selection procedure that eliminates molecules, which do not fulfil the criteria set by the experimenter.

3.4.3 The RNA model

In simple evolution experiments with RNA molecules genotype and phenotype are two different properties of one and the same molecule: The genotype is the sequence, the phenotype is the RNA structure with minimum free energy (mfe) and accordingly, unfolding of the phenotype is encapsulated in structure formation under thermodynamic control. Figure 3.20 presents a sketch of RNA structure formation, which occurs in two steps: (i) The single stranded molecules fold into a *secondary structure* through forming double helical stacking regions (shown in color) through base pair formation (Watson-Crick, $\mathbf{G} \equiv \mathbf{C}$ and $\mathbf{A} = \mathbf{U}$ as well as $\mathbf{G} - \mathbf{U}$ pairs),⁹ and (ii) the secondary structures fold into a full 3D structure on addition of two-valent cations [TLWK01]. RNA secondary structures have physical meaning as coarse grained versions of full structures.¹⁰ At the same time they are sufficiently simple combinatorial objects and allow for an equivalent representation as a string of chain lengths n over an alphabet with three symbols $\mathcal{C} = \{ (,) , \cdot \}$. The parentheses denote base pairs and fulfill the common rules of mathematics, and the dots symbolize unpaired bases. Accordingly, only a subset of all strings represents allowed structures, since (i) the number of left-hand parentheses has to match exactly the number of right-hand parentheses, and (ii) no parenthesis may be closed before it was opened (e.g., $(\cdot) \cdot \cdot (\cdot)$ is an invalid structure). The symbolic representation of secondary structures is particularly well suited for investigations on sequence-structure relations (See figure 3.20). Algorithms for fast computation

⁸Considering the chemical formulas of DNA and RNA, thymine (**T**) is a derivative of uracil (**U**); **U** compared to **T** differs only by the lack of a methyl group in position 5. The backbone of RNA carries an additional hydroxy group in position 2' of the ribose moiety.

⁹The number of lines between the two nucleotides is a measure for the strength of the interaction: **GC**>**AU**>**GU**.

¹⁰Coarse graining of RNA structures can also be interpreted in terms of free energy contributions since the formation of stacked base pairs is by far the largest stabilizing contribution to the total free energy of RNA structure formation.




Figure 3.20: **The structure of phenylalanyl-transfer RNA from yeast.** The sequence containing several modified nucleotides (**D,M,Y,T,P**) is shown on top. Under suitable conditions the sequence folds into a cloverleaf structure (drawing in the middle) called secondary structure that consists of rigid double helical parts and flexible loops and free ends. The secondary structure is equivalent to a string representing base pairs as parentheses and single nucleotides as dots. Changes in conditions, addition of $\text{Mg}^{2\oplus}$ cations, completes the rigid 3D-structure (For more details on RNA secondary structures see [Sch06]).

of RNA secondary structures are available [ZS81, Zuk89, HFS⁺94] (For an on the internet server performing RNA folding see, e.g., [Hof03]).

The relation between genotypes and phenotypes is visualized as a mapping from a metric space of genotypes onto a (metric) space of phenotypes (figure 3.21). The space of genotypes of chain length n over an nucleotide alphabet \mathcal{A} is denoted by $\mathcal{Q}_n^{(\mathcal{A})}$ and defined as follows: Every sequence of chain length n is a point in genotype or *sequence space*, and the distance between two sequences is given by the Hamming distance¹¹ $d_H(X_i, X_j)$. In the RNA model the phenotypes are the secondary structures, which are represented by points in phenotype or *shape space* \mathcal{S}_n . Since secondary structures are equivalent to strings over the three-letter alphabet \mathcal{C} it is straightforward to define a metric on shape space being, for example, the Hamming distance between the string representations of structures, $d_H(S_i, S_j)$.

The mapping from RNA sequence space onto RNA shape space can be written as

$$\begin{aligned} \Psi : \{ \mathcal{Q}_n^{(\mathcal{A})}, d_H(X_i, X_j) \} &\Rightarrow \{ \mathcal{S}_n, d_H(S_i, S_j) \} && \text{or} \\ S_k = \Psi(X_j), & j = 1, 2, \dots, n_k . \end{aligned} \quad (3.3)$$

By choosing the mfe criterion for RNA structures the assignment of structures to sequences is unique.¹² The number of sequences over the natural alphabet, $|\mathcal{Q}_n^{(\mathcal{A})}| = 4^n$, always exceeds the number of structures, $|\mathcal{S}_n| < 3^n$, where ‘smaller’ indicates the restriction on acceptable strings mentioned above. Accordingly, the mapping Ψ is many to one and not invertible. Sequences folding into identical structures are neutral with respect to the map Ψ and give rise to *neutral networks*. In the RNA model fitness is derived from structure as expressed by a mapping from shape space into the real numbers

$$f : \{ \mathcal{S}_n, d_H(S_i, S_j) \} \Rightarrow \mathbb{R}^1 \quad \text{or} \quad f_k = f(S_k) . \quad (3.4)$$

Although the relation between structure and fitness, encapsulated in $f(S_k)$, is not (yet) accessible by calculation, fitness values can be measured and obtained from plausible models. Different structures may have the same fitness value giving rise to a kind of neutrality in evolution that is superimposed upon the neutrality of sequence-structure map Ψ discussed before.

3.4.4 Modelling evolution

Simulation of RNA structure evolution provides detailed insights into the optimization process that cannot be obtained by the currently available experimental techniques. An individual simulation run monitors the search of a population of RNA molecules for a predefined target structure by replication, mutation, and selection. The computer experiments simulate a flow reactor (figure 3.22), in which the material consumed by the replication process is steadily replenished by an inflow of a stock solution. The population size is regulated by the outflow of reaction mixture that compensates the volume increase resulting from the inflow. The time development of the system is modelled by coupled chemical reactions. Conventional chemical reaction kinetics is based on ordinary differential equations (ODEs). Accordingly, conventional kinetics is unable to describe properly processes at low concentrations, since it neglects fluctuations that are important at small particle numbers. This is particularly important for evolutionary optimization, because every mutation

¹¹The Hamming distance between to end-to-end aligned sequences of equal length is the number of positions, in which the two sequences have different digits.

¹²The minimum free energies of RNA molecules will be almost always different provided the energy model has sufficiently high resolution. Exceptions are symmetric conformations with palindromic sequences, which are rare and which shall not be considered here.

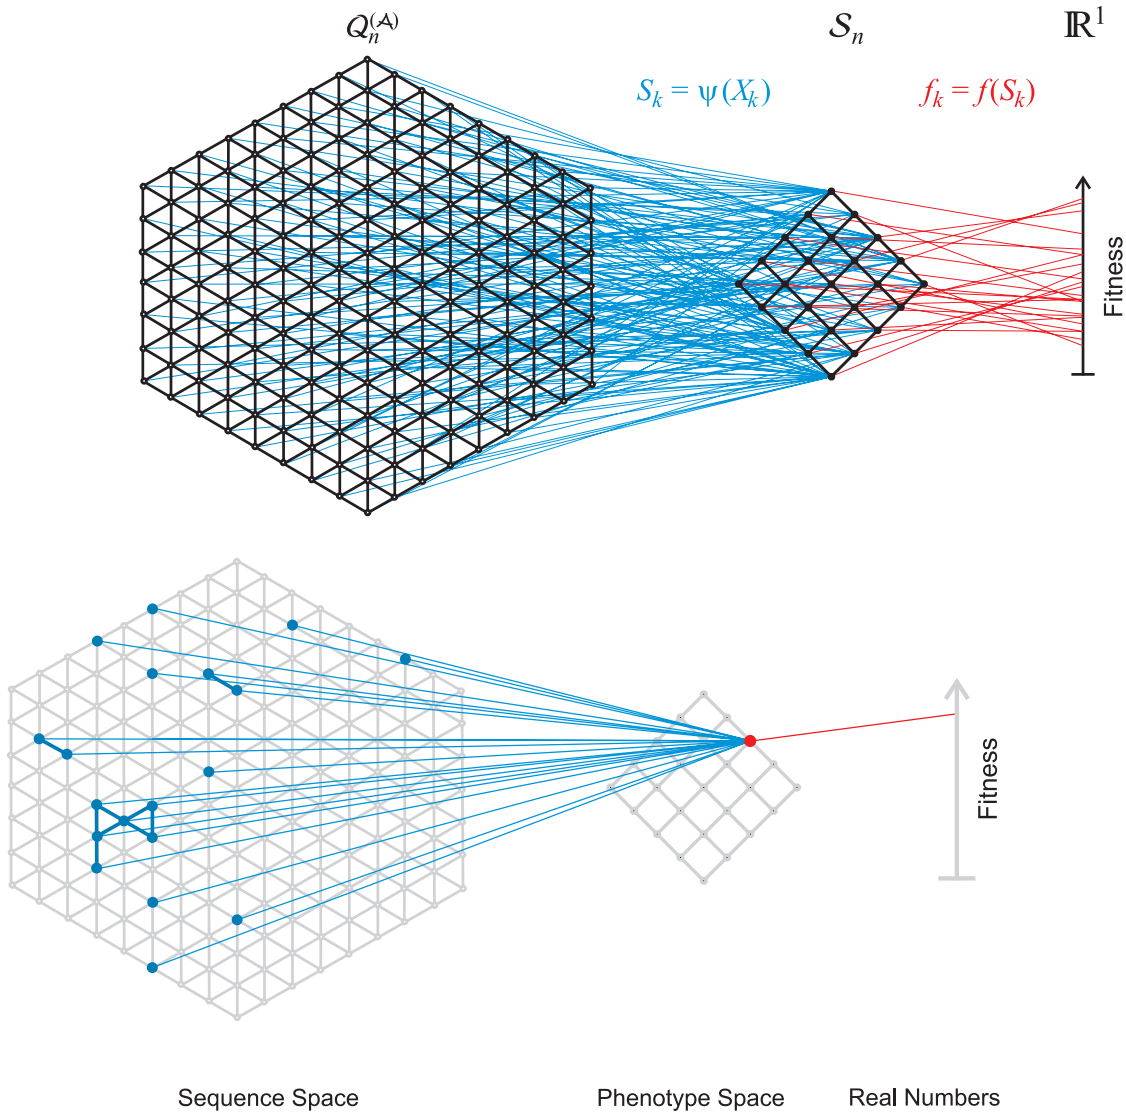


Figure 3.21: **Sketch of the maps from genotype space $Q_n^{(A)}$ onto a space of phenotypes S_n and further into real numbers.** Mapping sequences into structures of minimal free energy, Ψ as defined in equation (3.3), is many to one since for the natural nucleotide alphabets the number of sequences, $|Q_n^{(A)}| = \kappa^n = 4^n$, exceeds the number of structures, $|S_n| < 3^n$. Structures are evaluated by means of the function $f_k = f(S_k)$ (3.4). Here neutrality arises when $f(S_j) = f(S_k)$ holds for certain structures. The lower part of the figure highlights the part of the mapping concerning a single structure S_k . The preimage of the structure in sequence space is converted into its neutral network – shown in dark blue in the figure – by connecting all sequences of Hamming distance one.





Figure 3.22: **The flow reactor as a device for RNA structure optimization.** The reactor maintains off-equilibrium conditions by means of a flow, which supplies material for replication through inflow of stock solution. Molecules produced in excess are removed from the reaction mixture through an unspecific outflow. A population of N RNA molecules is subjected to replication and mutation. Minimum free energy secondary structures are computed and evaluated by means of the function $f(S_k)$ for all mutant sequences. The clover-leaf shaped yeast tRNA^{phe} (grey shape in the reactor) was chosen as target structure. Inputs of an evolution experiment *in silico* are the parameters (i) population size N , (ii) chain length n of the RNA molecules, (iii) the mutation rate p , and (iv) the initial population.

starts inevitably from a single copy. The model is therefore based on a stochastic process in many variables that is simulated by means of a straightforward algorithm developed in the nineteen seventies [Gil76, Gil77b, Gil77a]. In particular, the individual reactions are:



The indices 'i' and 'j' run over all molecular species in the reactor, f_i represents the individual replication rate parameters of the RNA molecule X_i , and the factors Q_{ii} and Q_{ij} are the relative

frequencies of correct replication events and mutations, respectively. In the simplest mutation model assuming uniform error rates p per site and replication the two frequency factors are: $Q_{ii} = (1 - p)^n$ and $Q_{ij} = (1 - p)^{n-d_H(X_i, X_j)} \cdot p^{d_H(X_i, X_j)}$ for single point mutations. The replication rate parameters are modelled such that the fitness increases with decreasing distance from target, for example $f_i = 1/(\alpha + d_H(S_i, S_T)/n)$ with S_T being the target structure and α an empirical parameter.

A simulation consists of building the average of a sufficiently large number of computed individual optimization runs called *trajectories*.¹³ A single trajectory is shown in figure 3.23. In this simulation a homogenous population consisting on $N = 3000$ molecules with the same random sequence and the corresponding structure (S_0 in figure 3.24) is chosen as initial condition. The target structure (S_{44} in figure 3.24) is the well-known secondary structure of phenylalanyl-transfer RNA (tRNA^{phe}) shown in figure 3.20. The mean distance to target of the population decreases in steps until the target is reached [FSS89, FS98a, Sch03]. Individual (short) adaptive phases are interrupted by long quasi-stationary epochs.

In order to reconstruct the optimization dynamics, a time ordered series of structures was determined that leads from an initial structure S_0 to the target structure S_T . This series, called the *relay series* [FS98b] a uniquely defined and uninterrupted sequence of shapes. It is retrieved through backtracking, that is in opposite direction from the final structure to the initial shape. The procedure starts by highlighting the final structure and traces it back during its uninterrupted presence in the flow reactor until the time of its first appearance. At this point we search for the parent shape from which it descended by mutation. Now we record time and structure, highlight the parent shape, and repeat the procedure. Recording further backwards yields a series of shapes and times of first appearance which ultimately ends in the initial population.¹⁴ Usage of the relay series and its theoretical background provides further insight into the evolutionary optimization process, allows to define nearness of phenotypes in evolution, and yields the basis for classification of transitions into minor and major events [FS98b, FS98a, SSWF01]. Figure 3.24 contains four selected structures of a relay series: Besides the start and the target structure, S_0 and S_{44} , we show S_9 being present for very long time in the relay series because there is no easy escape to a nearby structure with shorter distance to target, and S_{21} , which is one step before the formation of the cloverleaf through a major transition. Inspection of the relay series together with the sequence record on the quasi-stationary plateaus provides hints for the distinction of two scenarios:

(i) The structure is constant and we observe neutral evolution caused by neutrality of the map Ψ from sequences into structures. In particular, the numbers of neutral mutations accumulated on plateaus are proportional to the number of replications in the population, and the evolution of the population can be understood as a diffusion process on the corresponding neutral network [HSF96] (See also next subsection and figure 3.23).

(ii) The process during the stationary epoch involves several structures with identical replication rates and the relay series reflects a kind of random walk in the space of these neutral structures. This case corresponds to neutrality in the map f from shapes into fitness values.

Both classes of neutrality give rise to neutral evolution in the sense of Kimura's theory [Kim83].

The diffusion of the population on a neutral network or in a subspace of neutral structures is illustrated by the plot in the middle of figure 3.23 that shows the width of the population as a function of time [Sch03, GML⁺06]. The distribution of the population in sequence space broadens during a quasi-stationary epoch and then sharpens almost instantaneously after a sequence was

¹³In computer simulations different trajectories are obtained by repeating identical runs with different seeds of the random number generator. These seeds determine the sequence of stochastic events.

¹⁴It is important to stress two properties of relay series: (i) The same shape may appear two or more times in a given relay series series. Then, it was extinct between two consecutive appearances. (ii) A relay series is not a genealogy which is the full recording of parent-offspring relations a time-ordered series of genotypes.




Figure 3.23: **Variability in genotype space during punctuated evolution.** Shown is a single trajectory of a simulation of RNA optimization towards a tRNA^{phe} target with population size $n = 3000$ and mutation rate $p = 0.001$ per site and replication. The figure shows as functions of time: (i) the distance to target averaged over the whole population, $\overline{d_H(S_i, S_T)}(t)$ (black), (ii) the mean Hamming distance within the population, $\overline{d_P}(t)$ (blue, right ordinate), and (iii) the mean Hamming distance between the populations at time t and $t + \Delta t$, $\overline{d_C}(t, \Delta t)$ (green) with a time increment of $\Delta t = 8000$. The end of plateaus (vertical red lines) are characterized by a collapse in the width of the population and a peak in the migration velocity corresponding to a jump in sequence space. The arrow indicates a remarkably sharp peak of $d_C(t, 8000)$ around Hamming distance 10 at the end of the second long plateau ($t \approx 12.2 \times 10^6$ replications). In other words, every adaptive phase is accompanied by a drastic reduction in genetic diversity, $d_P(t)$. The diversity increases during quasi-stationary epochs. On the plateaus the center of the cloud migrates only at a speed of Hamming distance 0.125 per 1000 replications.




Figure 3.24: **Intermediate structures of structure optimization in the flow reactor.** We show the initial structure S_0 , structure S_9 , which is characterized by a particularly long plateau in the trajectory, structure S_{21} that is one step before the formation of the cloverleaf, and the target structure S_{44} .

created by mutation that allows for the start of a new adaptive phase in the optimization process. The scenario at the end of plateaus corresponds to a *bottle neck* of evolution. The lower part of the figure shows a plot of the migration rate or drift of the population center and confirms this interpretation: The drift is almost always very slow unless the population center ‘jumps’ from one point in sequence space to another point in sequence space where the sequence is located that initiates the new adaptive phase. A closer look at the figure reveals the coincidence of the three events: (i) beginning of a new adaptive phase, (ii) collapse-like narrowing of the population spread, and (iii) jump-like migration of the population center.

3.4.5 Neutral networks

The graph derived from the subset of sequences that are neutral with respect to the mapping Ψ through connection of all pairs of sequences with Hamming distance one is the neutral network. The neutral subset also called the pre-image of S in sequence space is defined by

$$\mathbf{G}[S] = \psi^{-1}(S) \doteq \{X | \psi(X) = S\} . \quad (3.6)$$

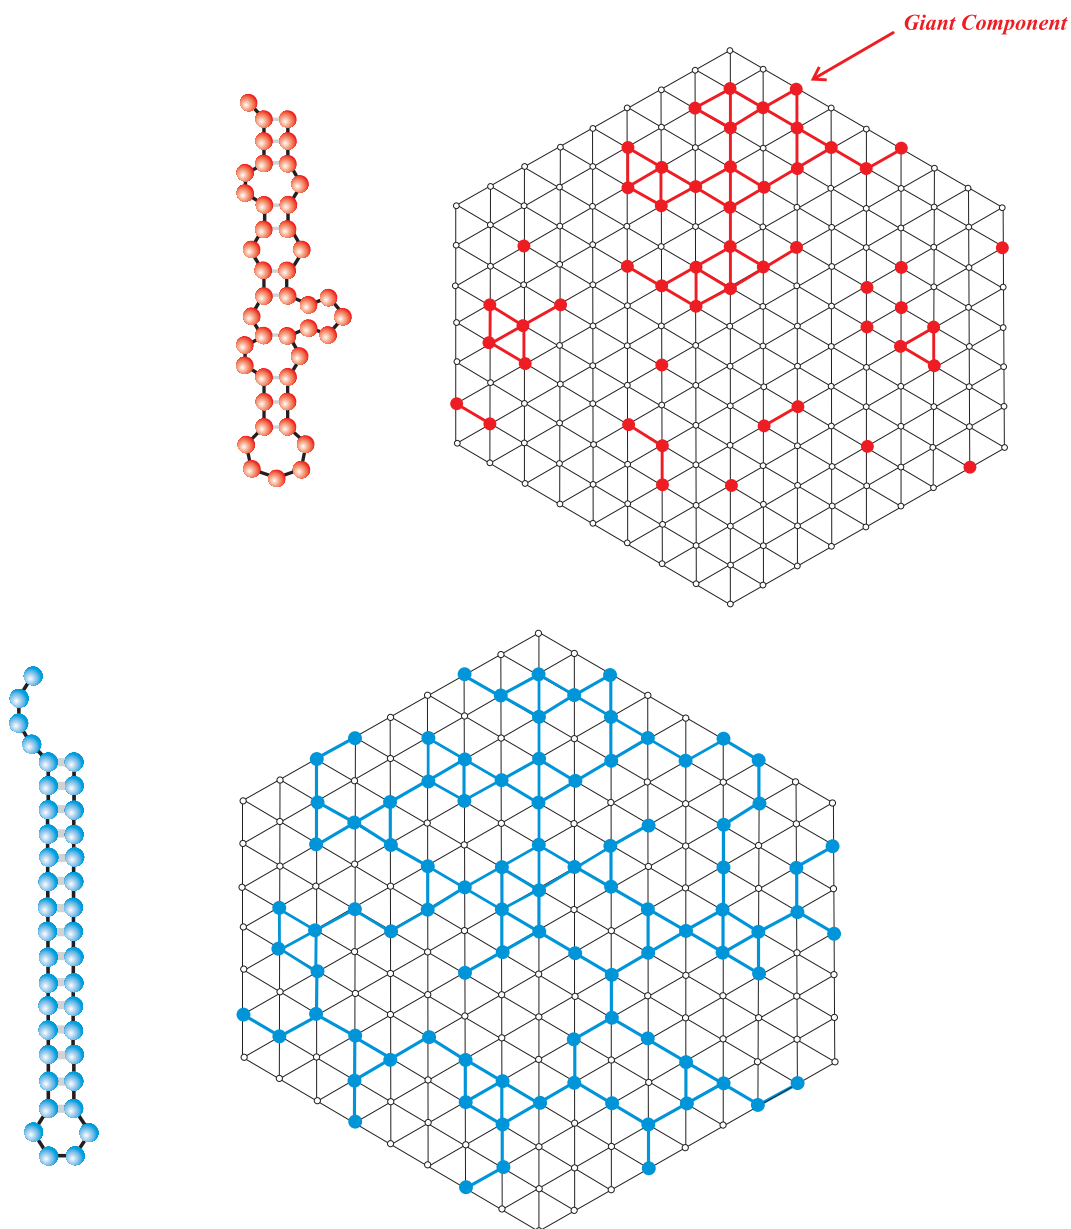


Figure 3.25: The sketch shows typical neutral networks of RNA structures. In the upper part we show a network that is characteristic for a low degree of neutrality. It consists of many components and, if random graph theory applies to the particular network, one component is much larger than the others ('giant component'). The lower part shows a network for a degree of neutrality above the connectivity threshold, which is fully connected.

Global properties of neutral networks may be derived by means of random graph theory [Bol85]. The most important property of neutral networks is connectedness (figure 3.25), which is determined by the degree of neutrality $\bar{\lambda}$. In particular, $\bar{\lambda}$ is defined and obtained by averaging the fraction of neutral Hamming distance one neighbors, $\lambda_X = n_{\text{nttr}}^{(1)} / (n \cdot (\kappa - 1))$ with $n_{\text{nttr}}^{(1)}$ being the number of neutral one-error neighbors and κ size of the nucleotide alphabet, over the whole network, $\mathbf{G}[S]$:

$$\bar{\lambda}[S] = \frac{1}{|\mathbf{G}(S)|} \sum_{X \in \mathbf{G}[S]} \lambda_X . \quad (3.7)$$

Connectedness of neutral networks is, among other properties, determined by the degree of neutrality [RSS97]:

$$\text{With probability one a network is } \begin{cases} \text{connected} & \text{if } \bar{\lambda} > \lambda_{\text{cr}} \\ \text{not connected} & \text{if } \bar{\lambda} < \lambda_{\text{cr}} \end{cases}, \quad (3.8)$$

where $\lambda_{\text{cr}} = 1 - \kappa^{-\frac{1}{\kappa-1}}$. Computations yield $\lambda_{\text{cr}} = 0.370$ for the critical value in the natural four letter alphabet. Random graph theory predicts a single largest component for non connected networks, i.e. networks below threshold, that is commonly called the ‘giant component’. Real neutral networks derived from RNA secondary structures may deviate from the prediction of random graph theory in the sense that they have two or four equally sized largest components. This deviation is readily explained by non-uniform distribution of the sequences belonging to $\mathbf{G}[S_k]$ in sequence space, which is caused by specific structural properties of S_k [GG⁺96b, GG⁺96a].

The existence of neutral networks is highly relevant for optimization on rugged landscapes¹⁵ As indicated in figure 3.26 adaptive walks minimizing cost function would soon end in a local minimum. They are unlikely to reach the globally lowest point of the landscape. Depending on population size smaller and narrow local peaks can be overcome by accumulation of mutants. Higher and broader peaks, however, are unsurmountable obstacles for the optimization process. Neutral networks change this situation drastically since the sequence corresponding to the lowest point in one dimension is commonly surrounded by neutral sequences corresponding to points of equal values of the cost function. As illustrated in the figure the population escapes from the trap by performing a random walk in an orthogonal dimension.

Connectedness of neutral networks allows for migration of the population over large portions of sequence space whereas neutral drift on disconnected random networks is confined to the individual components. Simulations of optimization runs with nonnatural nucleotide alphabets that do not sustain extensive neutral networks have shown that evolutionary optimization is indeed much more difficult on sequence spaces that do not support connected neutral networks (For example on the sequence space over the $\{\mathbf{G}, \mathbf{C}\}$ alphabet [Sch03, Sch06]).

In figure 3.27 we present a sketch of evolutionary optimization on sequences spaces with extended neutral networks. The population performs an adaptive downhill walk that is characterized by relatively fast decrease of the cost function until it reaches a point in sequence space from where no improvement is in reach by mutation. If this point is part of a neutral network the population proceeds by spreading on the neutral set (See the broadening of the population in figure 3.23). Neutral drift is slow compared to adaptive selection and takes place on a different timescale. In case the population reaches a sequence that has sequences with lower cost function in its neighborhood a new adaptive phase may begin, which leads to a network with lower cost function. The stepwise approach to target is continued until the population eventually reaches the target or it stays confined to a neutral network from which no further improvement is possible.

3.4.6 Multiple optimization criteria

Structure is only one property of the phenotype of an RNA molecule. Other fitness relevant properties are thermodynamic stability, stability against conformational change, stability against mutation or existence of a second long lived conformation.¹⁶ The goal of an actual optimization

¹⁵Rugged landscapes are characterized by large variation of the values of the cost function in the neighborhood of almost all points. Landscapes derived from biopolymer structures are typically rugged.

¹⁶Thermodynamic stability means low free energy of structure formation. Stability against conformational change is tantamount to a large free energy gap between the minimum free energy structure and the first suboptimal conformation. A structure is stable against mutation when the degree of neutrality, $\bar{\lambda}$ is large.

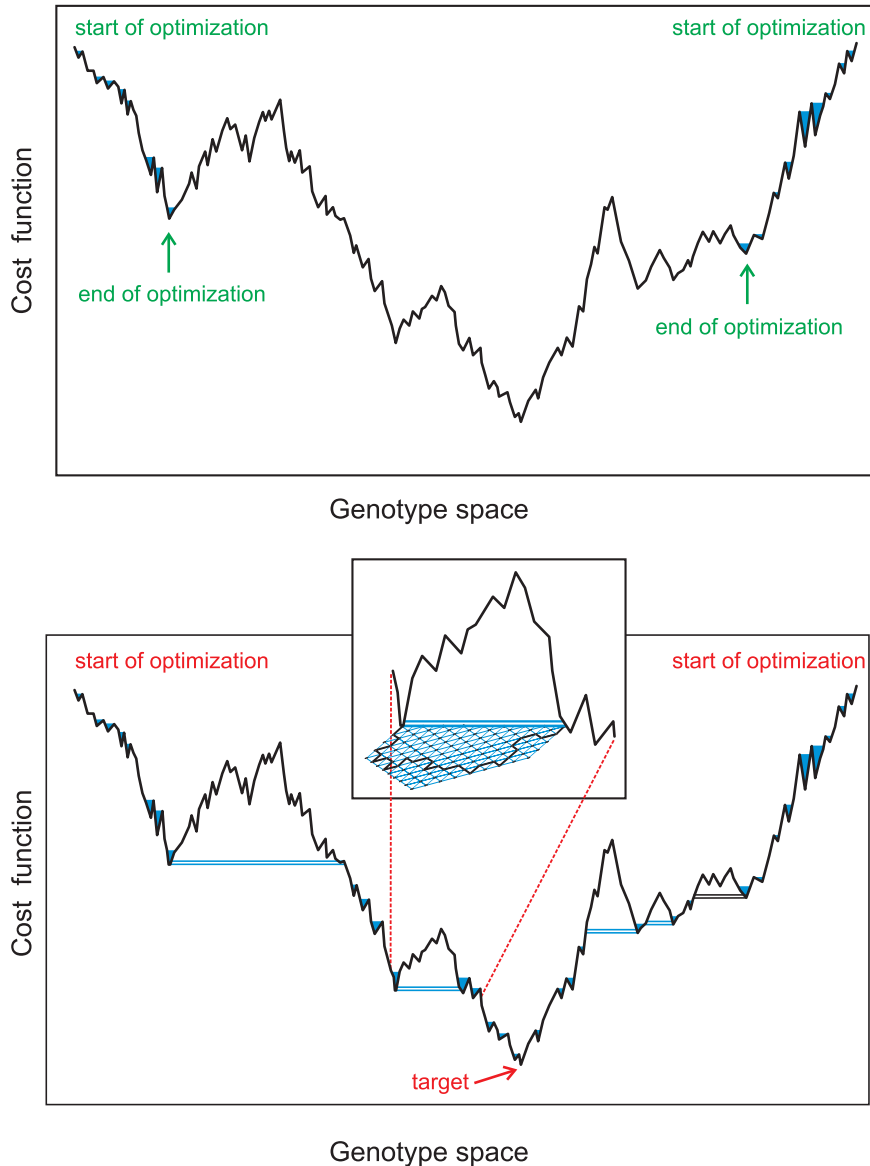


Figure 3.26: **The influence of neutral networks on optimization.** Adaptive walks minimizing a cost functions on rugged landscapes soon end at a minor locally lowest point. The population size determines the widths of clefts that can be bridged by accumulation of point mutations. Commonly, this width is rather small (upper picture). The lower part shows the influence of neutral networks that allow for ‘tunnelling’ beyond mountains through escape into another dimension where the network extends. The magnifying insert illustrates the escape into an orthogonal dimension and shows how a point that allows for further descent can be reached by a random walk on the neutral network.




Figure 3.27: **An optimization trajectory in RNA secondary structure space.** A target structure is approached in a stepwise procedure: Short periods of fast approach to the target are interrupted by long quasi-stationary phases during which the trajectory performs extended random walks on neutral networks. When a position is reached where a downhill path leads closer to the target, the trajectory leaves the neutral networks and approaches the next one. Such a position corresponds to a ‘tunnel’ in figure 3.26.

process may include several criteria for the molecule to be designed. Optimization according to multiple criteria is straightforward when we can assign importance to them in hierarchical order (figure 3.28).

In the design of catalytic RNA molecules with predefined functions it would make sense, for example, to assign highest priority to structure in the first step. Next, thermodynamic stability can be added as the second criterium for a search on the neutral net of the target structure and, eventually, one could introduce the requirement of stability against conformational change as the last property to be optimized. The result would be the most stable RNA sequence that forms a predefined structure with the largest possible gap between ground state and first suboptimal conformation.

Dropping or relaxing the hierarchical order of criteria leads to the well known problem in simultaneous optimization according to multiple criteria: There need not be a unique solution and one has to deal with a Pareto front or Pareto surface. For the properties of RNA molecules this problem has been addressed recently [SF03]. In reality, independence of optimization criteria will be the exception rather than the rule. Asking for the most stable molecule without the structural constraint will almost certainly yield a sequence that does not fold into the desired structure. In other words, molecules that form the right mfe structure will presumably have higher free energy than the most stable molecule. Analogous considerations show that many other pairs of properties do not allow for independent optimization. There is also a counterexample in the realm of RNA: Efficient or fast kinetic folding and thermodynamic stability seem to form an almost independent

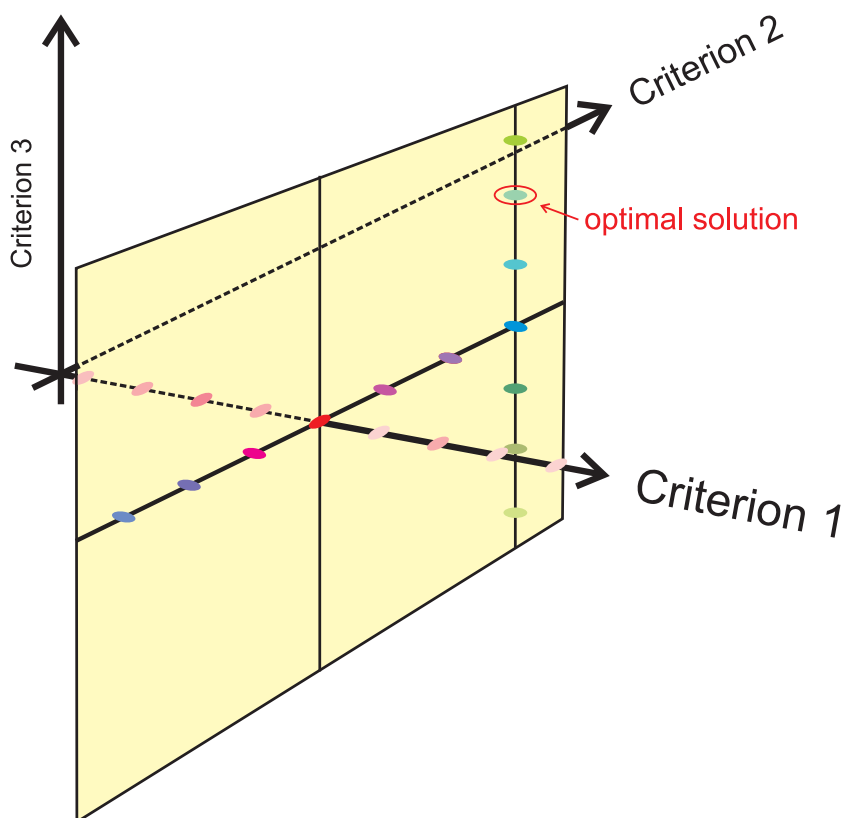


Figure 3.28: **Hierarchical optimization using neutral networks.** The sketch shows optimization according to three criteria. Criterion 1 defines a large neutral network that sets the stage for further optimization. Here we assume criterion 1 to be most important or to have the highest waiting factor in the cost function followed by the criteria 2 and 3. The procedure indicated leads to the global optimum if and only if the three criteria are independent. Otherwise the procedure has to be iterated and may lead to a Pareto front. An example from RNA structural biology is discussed in the text where the three criteria are secondary structure, thermodynamic stability, and uniqueness of minimal free energy conformation.

pair of properties [Sch06].

3.4.7 Applications of optimization on neutral networks

The occurrence of neutrality with respect to optimization is by no means restricted to a world of RNA molecules. Other biopolymers – proteins, for example, show very similar degeneracy with respect to the formation of coarse grained structures. Application of the concept of optimization on neutral sets of solutions are straightforward.

A simple example for neutrality based optimization following hierarchically ordered criteria is taken from a fictive hiring strategy. We assume a very large number of applicants from which one person should be chosen. The first criterion obviously is qualification and usually a relatively large subset of candidates fulfils this requirement equally well – we may consider these people as neutral with respect to this criterion. The typical second criterion is the performance in an interview. A third criterion could concern the personal skills of the applicant, which are useful for the work the person is thought to do. If still more candidates are left, a very last criterion might be avoidance of coincidence of names in order to be able to distinguish people more easily.

Other cases that start from a large number of coarse grained, equally suited initial solutions, which are split according to secondary or higher order criteria, are readily constructed. The major question concerns the usefulness of the neutral network approach. Indeed, we could optimize according to hierarchically ordered criteria without knowing about solution spaces and neutral networks. In case independent information is available on the nature and the structure of the corresponding neutral networks in sequence space, it is straightforward to derive optimal population sizes, hints on nature and optimal frequencies of variations. Point mutations are suitable on relatively dense neutral networks whereas insertions and deletions may be required in case of sparse networks with many components. Similarly, we would be able to answer the question whether recombination is useful or not.

3.5 Energy Engineering: Optimal Structures of Residential Heating Systems

Peter Roosen

Problem definition and scope

Developing innovative space heating concepts necessitates consideration of various factors influencing the pragmatic target function ‘warm, comfortable housing space’. Additional goals of the individual user must be considered, e.g. the accepted energy consumption, the costs conjoined with that, the flexibility of utilization, and many more. Besides the purely technical development of a respective heating system it is thus necessary to identify a suitable application niche relative to already established systems: It is not to be expected that a new design will outperform existing ones in every respect. There are too many application profiles and boundary conditions for the requested space heating demands. Prior to an extensive development of a novel solution the primary advantages need to be identified.

For every heating system there are several configuration parameters (like area proportions, insulation thicknesses, temperature niveaus, operation control concepts) exhibiting specific effects on the usually contradicting pragmatic target functions, like overall costs, cosyness primary energy consumption etc. . The targets are affected differently, depending on the foreseen usage scenarios (like constant usage of a flat vs. modulated usage of a utility building vs. stochastic usage of an event hall) clothing habits (strict business dress code vs. casual wear) and unchangeable parameters for a given building like window area ratio. Therefore the optimal system configuration has to be determined individually with respect to the potential relative target function ratios.

Varying the relative target function weighting and monitoring resulting major system parameter changes, like a switching of the basic heating system type, yields a notion on *decision stability*. A stable decision in this sense is a constant selection of one certain system type even in the light of slightly changing subjective target function weightings.

The method of choice in this paper for identifying dominant and stable ranges of a heating system is the pareto optimization. The set of optimal solutions for two or more partial targets, the pareto set, is determined, covering all possible relative weightings of the considered partial targets. Due to the fact that the individual target function value is calculated independently from any other and not traded against them there is no need to require an a-priori weighting. The targets are instead pursued simultaneously and independently.

In the scope of this exemplary study the following **target values**:

- a comfort measure
- the primary energy consumption

for the **simulated space heating variants**:

- gas fired floor heating
- gas fired wall radiator heating
- dynamic planar electric heating

are considered. The analysis is carried out for three **usage scenarios** and a floor area of 78 m²:

- Continuously used flat (living room, sleeping room, kitchen, hall) with differentiated usage profiles and respective heat gains in the different rooms

- Same flat with same relative usage profiles, but pure weekend usage (Friday from 18:00 h to Sunday 18:00 h)
- open-plan office with same ground area, 15 computer work places and flexitime usage (workstart in the interval 7:30 h to 8:30 h, combined with respective distributed work end in the afternoon).

The diversity of usages and the respective inner gains with their varying amounts and temporal distributions will show the differences in heating systems performance and their selective advantages for those tasks.

Simulation and Optimization

The aim of the reported work was to identify effects of specific usage scenarios and boundary conditions on the optimality of the target functions 'energy consumption' and 'comfort'. The desire to obtain significance for practical purposes led to the conclusion that the underlying system simulation should be as realistic as possible. Several energetic simulation systems, each cumulating several ten years of development times, are suitable for this purpose. Taking such a complete simulation package usually leads to the fact, though, that there is no possibility to interact with inner structures of the simulative calculations from the outside, i.e. the user interface. It is rather necessary to treat the simulation package as a whole as a kind of 'black-box' performer that may respond to parameterized simulation task definitions with its specific way of result response. The response, typically made available in the form of files, has to be scanned and parsed for significant target function information by the optimization procedure. A re-implementation of a simulation core is usually prohibitive due to the complexity of such a task.

The simulation system ENERGYPLUS (e^+) [Dep03b], provided by the US-american Department of Energy, seems especially suitable for optimizing the energetic and comfort behaviour of residential space heating systems. It combines a very great detailedness of simulated effects with a relatively simple system specification language in input files that may be artificially created and/or modified by a custom optimization module. A simulation task may be invoked by a command line call, referencing a respective pure ASCII input file of some 200 to 300 kB in size. The input file contains the individual settings of (a part of) the building to be simulated, like geometries, geographic orientation, wall structures, usage profiles, efficiency values for heating devices and such. They are defined by comma-separated lists depending in their structures on precursory keywords, relating to each other. Simulations are performed for real weather data sets defined with respect of temperature, humidity, and irradiation data for numerous cities of the world. Consequently solar irradiation gains through windows are considered in detail, with their geographic orientation being taken into account. We chose a weather data set of Düsseldorf/Germany for the optimization calculations documented below.

Some configuration parameters require a somewhat detailed intermediate treatment in transferring their settings from optimization to simulation module. As an example, the shift in an onset time of a heater needs to be reformulated into a complete rewritten 24 hours scheme in order to make e^+ recognize the change. For this task the macro preprocessor m4 [Sei00] was used. It creates the actual simulation input file with valid 24 h schemata by changing a template containing respective macros.

Upon request e^+ creates very detailed load profiles, energy flux reports etc., and also time-resolved comfort calculations according to the Fanger model [Dep03a] that are written into respective output files. In order to provide the optimization module with the necessary target function values for each previously defined system setting the output files have to be pre-processed in an appropriate fashion.

The Fanger Comfort Model and its Adaptation for Optimization Purposes

The Fanger comfort definition, being the most commonly adopted, is based on the six factors air temperature, humidity, air speed, mean radiant temperature (MRT), metabolic rate and clothing levels. (For an elaborate, in-depth discussion on Fanger's and other comfort models see [RdD97]) Fanger's model is based upon an energy analysis that takes into account all the modes of energy loss from the body, including: the convection and radiant heat loss from the outer surface of the clothing, the heat loss by water vapor diffusion through the skin, the heat loss by evaporation of sweat from the skin surface, the latent and dry respiration heat loss and the heat transfer from the skin to the outer surface of the clothing. The model assumes that the person is thermally at steady state with his environment. It was derived from experiments with American college age persons exposed to a uniform environment. The comfort equation establishes the relationship among the mentioned environment variables, clothing type and activity levels.

However, the equation only gives information on how to reach optimal thermal comfort by combining the variables involved. Therefore, it is not directly suitable to ascertain the thermal sensation (in terms of discomfort) of a certain person in an arbitrary climate where the variables do not satisfy the equation. Fanger used the heat balance equation to predict a value for the degree of sensation using his own experimental data and other published data for any combination of activity level, clothing value and the four thermal environmental parameters.

Even though the comfort model is based on objective thermodynamical interdependencies there is a substantial influence of subjectivity to the experimentally observed results (see also section 2.5.4). Accordingly test persons disagree to a limited extend on the experienced comfort, necessitating a statistical approach to represent their sensation as an averaged value. Therefore the Predicted Mean Vote (PMV) was defined that is expected to arise from averaging the thermal sensation vote of a large group of people in a given environment. The PMV is a complex mathematical expression involving activity, clothing and the four environmental parameters (see [Dep03a]).

The thermal sensation/discomfort is expressed as an index value, derived from the commonly used seven point psycho-physical ASHRAE scale, as summarized in Table 3.2. A PMV between -1 to 1 is regarded as the comfort zone.

Sensation	Hot	Warm	Slightly warm	Neutral	Slightly cool	Cool	Cold
ASHRAE scaling	1	2	3	4	5	6	7
Fanger PMV values	-3	-2	-1	0	1	2	3

Table 3.2: Thermal sensation PMV scale after Fanger, using a seven point psycho-physical ASHRAE scale

The interpretation of the Fanger PMV comfort value provided by e^+ needs some additional comments. In contrast to the discrete model originally devised by Fanger the comfort value returned by the simulator is a continuous one. But the general idea is preserved by modeling coldness discomfort notions as negative, hotness discomforts as positive values. In order to attribute comfort values to heating system configurations the hour-wise and room-resolved *absolute* Fanger values are averaged after weighting them with the number of inhabitants present in each one-hour period. Hence the goal of the optimization process is to strive for a positive bound value of zero as comfort measure. Unavoidable periods of overheating during summertime limit the lower bound of it to values well above zero, though, as the investigated spaces are only equipped with heating, not climatization devices. Underheating and overheating periods in unpopulated rooms do not contribute to discomfort as there is no person present to notice a discomfort.

The inhabitants' comfort assessment is closely coupled with their (simulated) clothing thickness. Changes in the range from light summer casual wear to thick winter pullovers yield significant differences in evaluation in otherwise identical thermal environment conditions. For practical purposes it is therefore not sensible to assume constant clothing conditions: It is definitely more realistic to assume the typical dynamical behaviour of the inhabitants to partially undress or add pieces of clothing within the respective socially controlled limits if slight notions of discomfort are experienced. Such an adaptive clothing schema is not feasible in e^+ , though. Here a fixed hourly schema for a set clothing thickness is to be provided.

To achieve a practically realistic statement a calculational trick has been implemented. The number of inhabitants of every room is halved, with one half being clad as light as possible and the other half as thick as possible. As the ideal comfort value is zero, a change in the sign of the Fanger value for the two halves is taken as an indicator that there exists a clothing thickness in the foreseen range providing maximum comfort. Accordingly, for that interval the practical comfort target function value is set to zero, designating the optimal one in the positive definite range of comfort assessment absolute values. If the signs of both halves are same, either the thinnest clothing is still too warm or the thickest is still not warm enough to obtain the maximum comfort. In that case the lowest absolute value will be taken as target value, assuming that the inhabitants will adopt their clothing to the nearest socially acceptable level. The latter changes significantly in context with the simulated scenario: The variation span for private living conditions is definitely larger than the one for the bureau business condition.

Optimizer adaptation

Since the simulator's internal organisation and the functional structure of the target function in its dependence on the available configuration parameters remain hidden to the user's scrutiny, the optimization module that is to cooperate with the simulator should react as insensitive as possible against non-differentiabilities and unsteadiness. Therefore a custom developed evolutionary optimization system was put to work [Roo99]. It provides a dynamic list population management, a Pareto set oriented multicriterial assessment schema and could quite simply be parallelized with respect to the invoked simulator runs by using the PVM system [PVM, GS92]. This simple parallel execution of up to seven simulation runs on separate TCP/IP network coupled computers was sufficient in the context of the given task as is shown in fig. 3.29. Without changing the



Figure 3.29: CPU load curves of four of the cluster of six computers working in parallel. The dead times of individual units appear statistically during the closure runs of each generation.

relatively inflexible generation paradigm the dead times imposed by idling processors at the end of a generation's simulation runs do not add up and reduce the overall calculational speed too much: The ratio $1/2 \cdot CPU\ count / individual\ per\ generation$, representing the statistical idle time per generation, is simply too small. The calculation time of one simulation run is typically about

40 sec (simulator machines' performance category: Celeron 1 GHz). Accordingly the speed loss due to necessary information transfer times was negligible for the less than 1 kB communication requirements per simulation invocation. Even when using more powerful processors the described concept will still be favourable and applicable without larger performance losses.

Heating system comparison for different usage scenarios

All scenarios were optimized for the same basic boundary conditions of floor area and geographic orientation. Furtheron it is assumed that identical usage profiles apply for adjacent units above and below the calculated one, so no net heat transfer is calculated via floor and ceiling. The residential units are simulated as four zone models (see fig. 3.30: living room, sleeping room, kitchen, hallway without heating). The bureau is rendered as one zone (open-plan office).




Figure 3.30: Structure of the residential units, both continuous usage and weekend use scenario. There is no heat exchange via floor, ceiling, and adiabatic walls to next flat on same floor.

According to the different usage profiles some hourly scheeas were adapted analogously: The forced venting by outside air ('infiltration') was significantly reduced in times of absence. The temperature schema was accordingly adapted for bureau (Mo – Fr) and weekend flat (Fr – Su) usage. The temperature settings for nightly drawdowns were applied during the other times.

Since the different heating systems exhibit very different heating dynamics — with the floor based heating being known as exceptionally slow — different pre-heating times have been introduced as optimizable parameters. So competitive comfort values can be reached, possibly conjoined with significantly increased energetic inputs, though.

The following configuration parameters were changed within sensible limits:

Technical section

Thickness of inner walls. This value describes the the thickness of the walls between adjacent rooms inside the calculation region. Since massive walls act as heat buffers this value has direct influence on the temperature characteristics resulting from short-term stochastic heat flux changes (solar irradiation through windows, ventilation events).

Thickness of wall insulation. The outer face of the facade is equipped with a respective heat insulation, interpreted as a polyurethane layer.

Maximum heating power. This value limits the energy output of the heating system, limiting the energy consumption as well influencing the comfort.

Heat transfer coefficient of the radiators. This value only exists for hydraulically operated heating systems and describes the heat transfer coefficients for the individually simulated rooms. In the case of floor heating the numerical value describes the total length of heating tubes which, for a specified width of the tubing, is a direct representative of the heat transfer coefficient.

Performance relating section

Temperature settings (regarded as constant heating control system targets over the whole simulated period)

Daytime temperatures in individual rooms. Temperature targets of living room and kitchen (together) and sleeping room (separate) are specified. The hallway is only passively heated by its energetic coupling to the other rooms. Living room/kitchen presets are interpreted as setting for the bureau one zone modelling.

Nighttime temperatures in individual rooms. Same specifications as for daytime temperatures, but given for the drawdown times. Same treatment in the one zone model, like with daytime temperatures.

Begin of daytime operation. This is, like the other temperature settings, interpreted as time of day at which the heating is switched to daytime operation; living room and kitchen.

Stop of daytime operation. Time of day when nighttime operation starts.

Preheating period. This period denominates the advance switch-on interval for the weekend usage simulation, to achieve a sufficient comfort level at the time of personal presences.

Absence temperature. Prevention of freezing and the required preheating period upon arrival is influenced by this temperature being set during times of absence.

So all scenarios contain at least 10 configuration parameters that are subject to optimization.

Continuously inhabited flat. The continuously inhabited flat accomodates a varying number of people. Following the assumption that the sleeping room is practically uninhabited during daytime, while the kitchen and the living room is not used during night time, there are periods of potential temperature lowering and hence energy consumption reduction without loss of comfort. Depending on the characteristics of the heating supply the rooms react differently if one tries to put this saving potential to work. The simulation periods contain two weeks per season.

First of all, fig. 3.31 shows clearly differing Pareto fronts for the individual systems. While floor heating and EDDH develop continuously bent curves the radiator heating seems to exhibit only piecewise steady ones that strongly jumps at a Fanger value of about 1. This jump coincides with respective error messages of the simulator that is no more comprehensible from the furtheron produced numerical results. The very small energy inputs observed in that range of the Pareto front seemingly lead to a non-converged mode of the simulation. Therefore the wall radiator results should only be considered at Fanger values less than one. This first qualitative result should be taken as a rather important issue in simulation-based numerical optimization: The optimization process itself may drive the simulator into parametric regions where it defects and delivers unreliable results without the tutoring human to notice it. So, in complex and unclear situations additional inspection facilities beyond the pure target function values should be provided to help the supervisor classify the conciseness of the calculated results.




Figure 3.31: Comparison of the Pareto fronts of three space heating systems with respect to comfort and primary energy consumption, under the condition of a continuous but differentiated usage of the available rooms. For discussion see text.

Further on the figure shows two separate wall radiator optimization run result sets to illuminate the statistical behaviour of results. Both runs yield the 'ideal' Pareto front in varying parts and stay inferior in others. Here the phenomenon of optimizational stagnation is evident: with evolutionary optimization there is no proof that the optimum really has been approximated. Repeated optimizations are mandatory, and even then there is no security on how far the real optimum frontier is still away, even if in most cases some asymptotic behaviour is likely.

In the continuously inhabited flat the unpopulated intervals in the different rooms seem to be too short to bring in some observable primary energy consumption effect from the rather rapid and dynamically operating EDDH, or the wall radiation system in a reduced amount, relative to slow floor heating appliance. Even though the EDDH end energy input (not shown in the diagram) is less than those of the two competing systems, the conversion losses in the production of the electrical energy (assumed effectivity of 40 %) leads to a primary energy related domination of the EDDH over the whole investigated range. The flat seems to be kept on a kind of quasi-static temperature level. The general, mean heat losses through the outside walls and windows are not evened out by inner gains and dominate the summed-up consumption picture.

Flat with weekend usage. The partitioning of the weekend flat is identical to the continuously inhabited one. During its times of usage the same inner gains are assumed, but only from Friday 18:00 h to Sunday 18:00 h. During the rest of the week all temperatures are set to the nighttime operation level, the forced ventilation is set to a very low rate, and no inner gains are assumed. The comfort rating is performed as usual: The absolute Fanger value is weighted with number of inhabitants. This leads in effect to a non-rating over the working-days part of the week. As the flat may severely chill out just before the weekend inhabitation starts again a freely choosable preheat time is provided to achieve sensible comfort values at the onset of each inhabitation period. This preheat time is a parameter in the performed optimization, leading to both increased cosiness and increased energy consumption and thus to a conflict in goals fulfillment.




Figure 3.32: Comparison of the Pareto fronts of three space heating systems with respect to comfort and primary energy consumption, under the condition of a weekend usage. See text for discussion.

Since the few usage periods in the simulated observation times of the continuously used room systems would presumably lead to erroneous results the complete simulation time was raised to 200 days per year, again assuming the weather reference data of Düsseldorf, Germany, provided by the e^+ system. This adaptation yields a number of comfort-assessed days comparable with the continuous usage situations. Nevertheless it is important to consider the intermediate, unpopulated days as well: Their settings influence both primary energy consumption and achievable comfort levels during inhabited periods.

In the weekend usage mode the electro-dynamical direct space heating concept can show its advantages: It shows a leading position in most parts of the Pareto tradeoff curve and surpasses the radiator solution in any case. Only in the high-comfort (but accordingly also high energy input) domain the floor heating concept is slightly superior. In the Pareto range around 1500 kWh/a most EDDH parameter sets show a very short preheat time prior to the 18 h arrival time of the inhabitants (17 h and later). This suffices due to the small thermic masses that need to be warmed up to obtain reasonable comfort values.

Bureau situation. Relatively large inner gains are observed in the bureau situation with the same floor area as the residential usage: An occupation density of 15 computer workplaces with the system switched on during normal bureau times is assumed. Accordingly a remarkable (dis-)comfort baseline is observed in the Pareto sets of all heating concepts (fig. 3.33). The heat produced by inner gains cannot be eliminated during summer days since no climatization, but only ventilation is assumed and simulated to preserve a certain level of comparativity. In addition, in the bureau environment the clothing order is somewhat more restricted with respect to very light clothing, leading also to a raised level of discomfort during hot summer days.

The relative positioning of the heating systems is somewhat similar to the results for the weekend used flat: While the floor heating concept dominates the range of high comfort (and high energy consumption) the low energy input (and combined with that, the low comfort) range is the domain of the EDDH. The radiator wall heating concepts exhibits an interesting distribution:




Figure 3.33: Comparison of the Pareto fronts of three space heating systems with respect to comfort and primary energy consumption, under the condition of a bureau usage. See text for discussion.

Its Pareto set mostly coincides with the dominated parts of the other two concepts, without ever being a leader itself.

Interpretation of results

The results allow some insights into the potential, but also the limitations, of space heating system optimization. First of all, the anticipated outcome is confirmed. 'There is no such thing as a free lunch' — or in other words: For each system in question one has to balance comfort requirements with an appropriate (primary) energetical input, although the amount of the latter may be systematically different across heat provision system types. Each heating system produces its own Pareto front in the comfort vs. energy input plots.

There is no general leader system completely dominating other existing concepts. Depending on the structure of the heating request — mainly continuous, intermediate with relatively long on/off periods, or short periodic changes as three example situations — the advantages and disadvantages of the supply systems become visible.

In a continuously populated residence the determining factor for primary energy consumption is the heat loss through walls and windows that has to be balanced by almost steady heat input during phases of cold weather. Here the overall conversion efficiency of the heat provision system plays the most important role for the primary energy input calculation. Since the production of the very high exergy carrier medium 'electricity' only coincides with relatively high conversion losses, this in turn leads to a general energetical inferiority of ohmic heat generation when same levels of comfort are to be reached. In the continuously used flat there are no major dynamical effects the highly adaptable electric heat production can utilize to lessen this general disadvantage.

The weekend use of a flat is almost a complete antithesis. Long chillout periods during the absence of the residents lead to the necessity to heat up relatively large volumes of thermally capacious, heat absorbing material. Here the established systems of radiators or floor heating need long preheat times in order to get the living space itself sufficiently cozy. After the inhabitants once again leave the flat, the heat energy mainly introduced into the building's walls slowly dissipates

into the environment without any further benefit. Contrasting to this behaviour the EDDH system mainly heats the space while leaving the surrounding walls mainly in their cold states during the rather short usage periods. So the total input heat energy per weekend is significantly smaller than with the other systems. This effect overbalances the conversion losses.

The slow start profile of the bureau scenario, due to the slow crowding of the workplace according to the simulated flexitime usage, once again seems to promote the relatively slow floor heating system, at least for sufficiently high comfort levels. The dynamic effect that would presumably promote the EDDH seems not large enough to overcompensate the general electrical conversion loss on the primary energy target function.

Although these results are quite convincing and can well be interpreted in terms of general scientific concepts, they rely strongly on the modeling input flexibility of providing the simulation model with practically relevant detailedness — and therefore on the effort in setting up the model. Some of those issues shall be scrutinized in more detail in the next paragraphs.

One major aspect of an EDDH application is the *very* fast response of the comfort value on any control action performed to the actual heat input into the room. This response is in the order of minutes, while the simulation system will only yield reasonable comfort values on the basis of 10 minute steps. A faint approximation is the assumption that each change in population of the simulated space, and any change of inner gains, just coincides with full hours for which the time schedules must be defined for the e^+ simulator. A more detailed simulation would require both many more time steps and a simulation system explicitly taking into account dynamic heat distribution processes in the surrounding walls. This in turn would multiply the required simulation time which was already at its limit for the available hardware.

e^+ provides some automatisms for adapting the actually input heating energy into the room, as well as for the venting in case of overheating etc.. Although these procedures are already quite elaborate with respect to realistic scenario modeling, it does not suffice to describe pragmatic personal actions. Just to give some examples, what is *not* sufficiently described: i) When people populate a flat the doors between rooms tend to be left open in case of rather similar temperature distributions while they are usually closed more rapidly if temperature gradients occur. ii) The comfort feeling in a surrounding with a temperature gradient — both in time and space — is not covered, as the Fanger model does not support this boundary condition. Even if people feel slightly uncomfortable due to a slightly too cold surrounding this feeling is somewhat compensated if the surrounding tends to get warmer. iii) If there are different temperatures in different rooms a person wanders frequently between the sensation will be different compared to his stay in just one of them. If a room is used only for short times (like some minutes in each case), but this rather often, a discomfort value attributed to this room will be in practice less perceived in comparison to a person staying in that environment for a longer time.

There is one additional aspect that cannot be covered even by the best possible situation-aware modeling: The potentially changing room assignment in a given complete space. Exchanging a bedroom against a living room, or splitting a large living room into smaller subsections that become somewhat autonomous in their usages, changing cooking behaviour — like switching from conventional cooking with large inner heat gains to microwave ovens with almost no external heat production, or exchange of slow bureau computers to faster but more energy demanding ones: Those usage changes cannot be foreseen, so another pragmatic optimization target should be kept in mind as very important for most cases: the usage flexibility.

Trying to integrate this aspect leads to the frontier of modeling complexity and the possibility to interpret obtained results, as there is no objective method in sight presently to economically (with respect to modeling setup) take this into consideration. The only resort that may address this problem might be the qualitative comparison of the magnitude of calculated scenario differences

and their attribution to the probability of happening scenario switches.

3.6 Civil Engineering: Structural Design and Process Shaping in Consideration of Uncertainty Phenomena

Dietrich Hartmann

Introduction

Structural design of contemporary systems in civil engineering, as well as the shaping of the processes associated with such systems, is characterized by an increasing complexity. This constitutes a big challenge to engineers which can only be mastered by means of rigorous computer-aided approaches. Customarily, engineers also attempt, to figure out what are the best solutions for their systems and processes. It comes, therefore, not as a surprise that, to obtain practicable results, the computer models established must to represent largely unaltered real world scenarios. Needless to say, that this general demand holds for many other scientific disciplines beside engineering. Thus, working solution concepts can be analogously transformed to other domains.

Due to the complexity, however, computationally intensive simulations are indispensable, with regard to the number of simulations and for the purpose of verification and validation. Not until that, a systematic evaluation of the simulations may allow for optimized systems or processes. To this end, the optimization or simply the amelioration of initial proposals is measured with respect to desired objectives, serving as a quality criteria subject to prescribed constraints. The high complexity additionally calls for multidisciplinary solutions and necessitates multi-level, multi-scale and even multi-paradigm models, leading to micro-, meso-, macro- and/or super-models interwoven to each other.

In the following, three different examples are dealt with which are to demonstrate how the handling of uncertainty in complex structural engineering problems can be accomplished, according to the general discussion in section 2.4.4, on page 33 ff. The first example is concerned with the lifespan-oriented design of large scale structures subject to time-variant stochastic loading, taking into account deteriorations. In the second example, the computer-aided destruction and collapse of complex structural systems using controlled explosives is discussed with respect to the realization of fuzzy randomness. Finally, the third example addresses the structural optimization taking into account stochastic imperfection in the geometry induced during the erection process.

Lifespan-oriented design

The lifespan-oriented design of structural systems in civil engineering is in the focus of a Collaborative Research Center (SFB 398), funded by the German Research Association (DFG). In this center, an own specific research project among others is dealing with the optimum design of steel structures accounting for the deteriorations and damages induced during the utilization of a structural system. This is to say that the design problem is transformed into an equivalent structural optimization problem.

For steel structures, in particular, the design against failures due to fatigue phenomena often becomes crucial. To give an example, industry halls composed of frames and used as a first reference system may fail under repetitive, stochastically multiply correlated as well as in stationary wind actions.

A further example are steel arched bridges (Langer's beams) which are considered as the second type of a reference system.

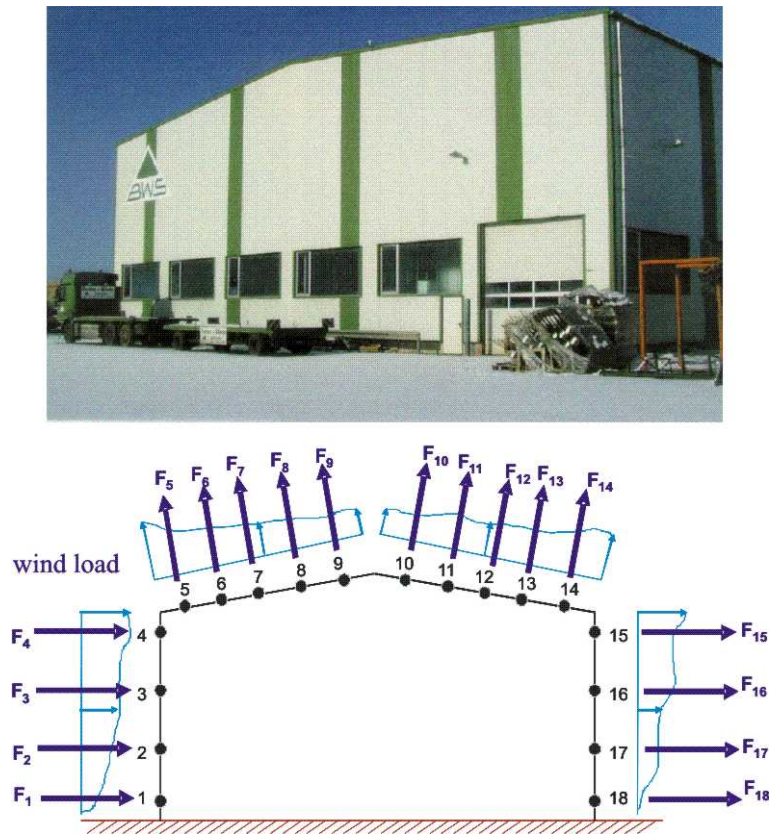


Figure 3.34: Industry hall subjected to wind loading



Figure 3.35: Steel arched bridges

Surprisingly, it has been assessed that even newly erected bridges show failures in the hanger connection plates linking the vertical arch hangers with the main girder of the bridge.



Figure 3.36: Hanger connection plate

Extensive research and testing have figured out that mostly vortex induced across vibrations cause the damage (cracks in the plate), not rarely shortly after the erection. Again, the randomness of the wind actions, even if stationarity is prevailing, affects the damage-sensitive vibrations within a so-called lock-in domain.




Figure 3.37: Vortex-induced vibrations (lock-in effects)

From these remarks it can be concluded that the lifespan-oriented structural analysis and design is substantially governed through a plenitude of uncertainties in the stochastic load processes

(actions) with regard to appearance, duration and intensities. Accordingly, the structural response in terms of displacements, vibrations, strains and stresses have to be regarded as time-variant stochastic processes themselves yielding stochastic deteriorations and a degradation in the total system. Of course, the relevant structural data have also to be treated stochastically in terms of basic variables (simple stochastic variables).

To elucidate, in principle, how randomness as one of the categories of uncertainty is incorporated into engineering design the afore mentioned industry halls are contemplated, again (Fig. 3.34). In this case a multi-level approach is applied to appropriately capture the real world behavior of the interactions between loading, structural response and deteriorations. By that, the time-variant computation of the deteriorations and the damage is accomplished by two consecutive analysis models (i) a first-model (Fig. 3.34) is used by which the stresses due to extreme turbulent wind velocities in a micro-time scale ($T \approx 10$ minutes) are computed (ii) a second model is taken to estimate the stresses due to wind lading in the macro-time scale (Fig. 3.38).




Figure 3.38: Intermittent Continuous Stochastic Process

Based upon this micro-macro separation, the fluctuations in the micro-model can be assumed as piecewise stationary and Gauss-distributed while the actions over the long term are represented as a stochastic intermittent continuous pulse process (Markov renewal process). The deteriorations induced in the micro-time scale are then accumulated, according to the pulse process and using stochastically a modified linear Palmgren/Miner (S/M)-rule. By means of the above described damage accumulation, representing a first passage problem, the failure probability P_f due to fatigue can be assessed. The value for P_f is not allowed to exceed a given admissible limit P_{adm} . This requirement, then, forms a stochastically non-linear and time-dependent constraint for the structural optimization problem besides side constraints for the selected optimization/design variables. For the steel frames considered here, according to (Fig. 3.38) only two sizing variables are defined: The height X_1 and the width X_2 .

The objective function is the cross-sectional area $A(X_1, X_2)$. Thus we have the following optimization problem

The solution of the exemplary design problem can be learned from (Fig. 3.41) where the optimization domain is depicted in the X_1/X_2 -space.

As illustrated, the stochastic constraints for the failure probability scatters around expectation values (see bell-shaped curves in (Fig. 3.41) and creates two segregated sub-domains: (i) a first domain containing green points (no failures) and a second domain of red points representing failures. The optimum is a set of points (encircled in blue) the size of which depends on the selected critical fractile.




Figure 3.39: Optimization model for framed system (industry hall)

$$\min \left\{ \begin{array}{l} Area(X_1, X_2) \\ g(\mathbf{X}, \mathbf{Y}, \mathbf{Z}(t)) = P_f(\mathbf{X}, \mathbf{Y}, \mathbf{Z}(t)) - P_{adm} \leq 0 \end{array} \right. \left. \begin{array}{l} 0,1 \leq X_1 \leq 1,0 \\ 0,1 \leq X_2 \leq 0,3 \end{array} \right\}$$

Figure 3.40:

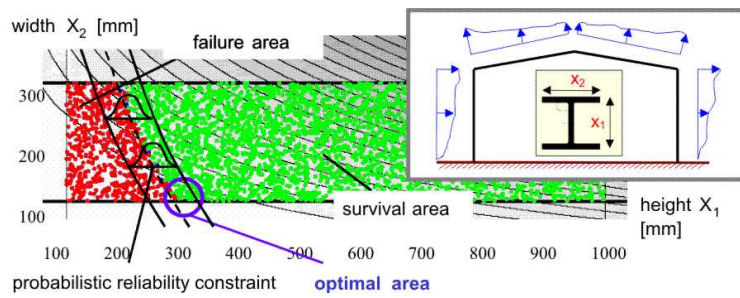


Figure 3.41: Solution domain for stochastic optimization

Computer-based demolition of complex structural systems

In many suburban densely populated areas there are existing buildings and constructions that arrive the expiration of their service life because of insufficient structural quality, changed requirements regarding utilization or simply unacceptable layout. As a consequence, an increasing number of buildings (departments store, administration buildings, multi-storey buildings etc.) have soon to be demolished. With respect to emerging costs, time and obstructions of people, traffic or inner-city infrastructures the demolition by using controlled explosives is mostly the best technique to tear down structures. This, however, only applies if the collapse, triggered through the explosive charges, elapses according to schedule and does not cause the slightest collateral damages, neither for human beings nor in the adjacent infrastructure. The following figure



Figure 3.42: Instances of explosions in the practise

exemplifies characteristic blast scenarios that took place in intercity areas such as the explosions for the demolition of (i) a library in Dortmund, (ii) a high-rise residential building in Hamburg and (iii) an administration building in Wuppertal. All of these demolitions have been attended and analysed for research purposes with the objective of creating a computer-based simulation system by which the real world collapse due the explosion can be mimiced. Based on this simulation, an optimal demolition strategy is

to be sought in such a way that the discrepancies between the provoked debris pile and a prescribed collision domain is "minimal". The research is carried out in the Research unit "Computer-based destruction of buildings" that, again, is supported by funds of the DFG, linking together Structural Mechanics, Computational Mechanics, Structural Concrete and Engineering Informatics.



Figure 3.43:

Everybody who has watched real world blast demolitions of complex buildings on the TV, or as spectator, can imagine that a plethora of imponderabilities may prevent the desired collapse result. Such imponderabilities may be caused by both data uncertainties (structural data, explosion parameters e.g. regarding volumes of charges, time sequences of ignition, efficiency of explosive charges) as well as model uncertainties (range of structural behaviors, contact and impact phenomena during the collapse, wind actions, etc.). If all relevant aspects are to be modeled realistically, then, the total aforementioned subcategories for uncertainty modeling have to be materialized, i. e.

- randomness
- fuzziness and
- fuzzy Randomness.

The application of randomness is possible if stochastic variables (basic variables) can be identified, if a sufficiently large universe is available and the mathematical principles of randomness are valid. Material properties, like the modulus of elasticity, Young's module, etc. are good examples for that. By contrast, fuzziness is more suitable if qualitative issues have to be captured, based upon subjectively defined membership functions for uncertain quantities. Herewith, a new type of interval arithmetic can be introduced instead of using conventional numerical concepts. Finally, if items or parts of the stochastic quantities (basic variables or random processes) contain incomplete information or violate the principles of pure randomness, then, fuzzy randomness fills the gap between objective randomness and subjective fuzziness within the same common body of representation. Beside the representation of uncertainty a further fundamental solution concept is mandatory. Due to the complexity of the simulation-guided optimization problem, the original problem has to be decomposed and distributed into subproblems, representing different problem levels with respect to space and time. In the Research Unit project, four distinct interacting levels are recognized to map the real world behavior of the collapse cascade appropriately. The subsequent figure (Fig. 3.44) illustrates the spatio-temporal multi-level approach more detailed.

Some few elucidations: As can be seen, on the micro-level the local effects of the explosive charges are modeled using a smooth particle hydrodynamics code and a net-free Galerkin method (hereby time scale ranging from micro- to milliseconds while space dimensions [L] are between mm and cm). In a micro-/meso-level sub-model the shock waves and vibrations are computed based on a finite element method associated with explicit time integration and element erosion ($\Delta t \approx$ milliseconds, [L] between cm and m). On the meso-/macro-level of the structural system an inventory control determines the random character of the structural data (concrete parameters, reinforcement assembly, etc.). Also, three dimensional strain-stress characteristics at critical sections are established which are employed in the collapse cascade of the total system level. The domain of influence is in meters. On the total system level the collapse cascade is simulated, which spans from a few hundredth to about hundred seconds, and comprehends the total structure and its fragments after the explosion. The complexity of the collapse cascade requires a flexible




Figure 3.44: Spatiotemporal in blast simulation

approximation model where a multi-body approach fits best, taking into account the essence of the three preceding sub-models. According to the multi-body model, bodies are regarded as objects interacting with other objects: Thus, an object-oriented implementation is accomplished that also allows for a mapping of the event sequences during collapse.

From the graphical diagram in Fig. 3.44 it can be recognized, that on all four sub-levels fuzzy randomness is integrated to embed the relevant uncertainties as already described above. For a prototypical system a high-rise concrete skeletal building, which covers the most significant collapse mechanisms, the first simulation results of the chosen fuzzy-random multi-level concept are available. In Fig. 3.45 the collapse of the building after the explosion is visualized. Also, the debris hill is portrayed.




Figure 3.45: Collapse simulation using fuzziness

Within the test simulation simplistically two quantities, the rigidity of the structure and the friction coefficient, have been fuzzificated in terms of two triangular membership functions. This approach yields an uncertain range for the debris radius (Fig. 3.46) where the α -value represents the level of membership: The value $\alpha=1$ means full membership while $\alpha=0$ (left and right of $\alpha=1$), marks vanishing membership.

Imperfection-sensitive structural optimization

Imperfections in a structural system, induced through construction and erection, are a natural phenomenon that never ever can be totally excluded. In general, as geometric imperfections may occur. Here, the focus is on geometric imperfections which lead to more or less aberrations from the planned geometric shape of a structure. Geometric imperfections can become dangerous in




Figure 3.46: Uncertain debris radius

structures subjected mainly to compressive stresses because sudden collapse can happen. Typical compression-subjected structures are shown in Fig. 3.47.




Figure 3.47: Examples for compression-subjected structures

In particular, if it is attempted to provide an optimum design for the aforementioned structures, using a structural optimization model, the optimum solution computed may be counterproductive. The reason for that is that the "optimal" shape obtained by numerical optimization may show an extremely sensitive behavior and fail due to the already smallest geometric imperfections. To explicitly exclude such catastrophic failures with a reasonable degree of probability, the uncertainties due to geometric imperfections have to be scrutinized and incorporated into the shape optimization model. This model has been created in the dissertation project of M. BAITSCH [BAIT03] at the author's institute where again a multi-level approach is required. For the reference structure displayed in Fig. 3.47, (left) the relevant uncertainties due to geometric imperfections have been successfully represented by means of stochastic fields, in association with a property chosen pdf and correlation functions. Starting with the prescribed perfect geometric shape of the structure the introduction of random variables (imperfection variables) for the stochastic fields allow the

assessment of the worst possible imperfect geometric and, accordingly, the corresponding most inappropriate structural response. To find out the unknown worst imperfect shape it is assumed that the amplitudes against the perfect geometry cannot be arbitrarily large. In harmony to that fact, the q -dimensional envelope (ball) in the space of the uncorrelated imperfection variables is determined. Hereby, the characteristic length (radius) of the envelope represents a prescribed fractile being a measure for the gradual assessment of uncertainty. Based on this concept, on a first level of the total optimization model the worst possible imperfect shape is computed by a maximum optimization (see Fig. 3.48).

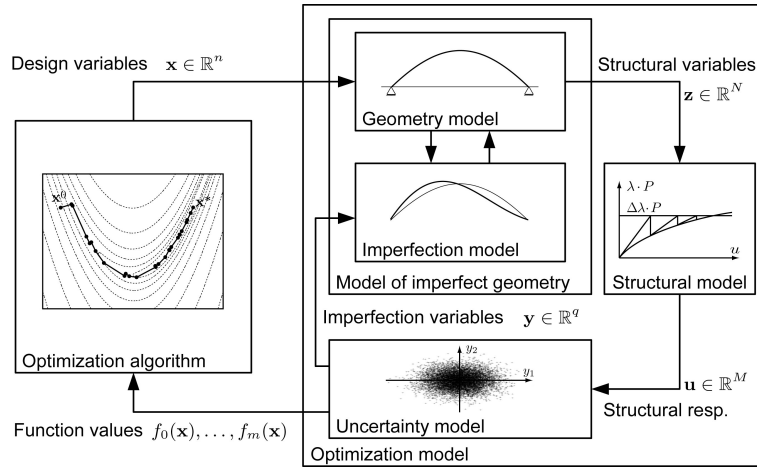


Figure 3.48: Two level optimization concept

Having determined the worst imperfect geometric a conventional structural optimization is carried out, on the second level (see Fig. 3.48). As illustrated, the uncertainty model yields the imperfection model for which a geometrically non-linear finite element analysis is performed. This analysis forms the kernel for the structural optimization in particular for the constraints of the optimization iteration. The design variable are mode coordinates of the arched girder. As objective criterion the strain energy is taken by which the robustness of the structure against sudden failure can be ensured. In Fig. 3.49 the optimization history is mirrored.




Figure 3.49: Optimization results

The picture at the top of Fig. 3.49 demonstrates the cascade-shaped reduction of the initial solution where the toggling between the anti-optimization and the structural optimization takes place. The bottom charts of Fig. 3.49 showing the initial and optimized structure, the corresponding displacements as well as the governing imperfect shape.

Conclusions

The actual endeavors in scientific structural engineering to understand highly complex systems and processes, realistically, definitely necessitates the introduction of different categories of uncertainty as well the integration of subjective assessments described in a mathematized fashion. By that, conventional numerical methods deployed in Computational Mechanics already for a long time can be linked to methods stemming from the Probability Theory and Computational Intelligence. The consequence is that such an enhancement of Computational Engineering opens new directions to cope with what has been denoted in Cognitive Science as "restricted rationality". In the above chapters it could be demonstrated that there are effective and sufficiently accurate ways and means to capture the sophisticated non-deterministic behavior of selected structural engineering, prudentially. Three applications examples have been scrutinized in which a wide band width of computer models for the acquisition of uncertainty phenomena could be pointed out: (i) the lifespan-oriented design of structures, (ii) computer-based demolition of complex buildings and (iii) the imperfection-sensitive structural optimization.

The essential benefit in all three cases is the sustainable improvement of the safety of the systems used and the processes activated. Hence, the vulnerability of the systems and the endangerment of human lives can be drastically scaled down. In this context, it is significant that the improvements or reductions can not only evaluated qualitatively but also quantitatively. That is to say uncertainty can be better gauged and better detected than before, in advance to hazardous events.

3.7 Logistics: The Complexity of Operational Transport Optimisation (Kopfer)

Describe the multitude of boundary conditions observed in real-world freight delivery problems (time windows, packing/stacking conditions, ...) and the methods that have been established to treat them. Elaborate on the nature of respective heuristical approaches and their limitations. Mention uncertainty aspects, like stop-ups and other casuals during delivery tours, and their effect on tour optimalities.

3.8 Innovation management: Assisting the development of optimal innovation processes

Martin Möhrle

Introduction

Today, the planning of innovation processes is supported by a large number of instruments, many of which are specifically designed for innovation management purposes, such as innovation evaluation, cost/benefit analysis, and lead user analysis (see [BCW04] [SG02] and [vH05] for overviews). Many other instruments come from classical project management, e.g. project structure planning, deterministic network planning, stakeholder analysis, and project controlling [Ker05]. These project management instruments are regularly incorporated in software tools like Microsoft Project and others.

Although the aforementioned instruments have indeed proved to be helpful for quite a few innovations, some innovations bear specific characteristics, which lead to a need for more advanced instruments and modeling. The characteristics in question are (i) the multiple risks and uncertainties, by which these particular innovations are influenced and which may be traced back to different origins, such as market demand, technological progress or administrative regulations. Also, there are (ii) multiple goals situations underlying those innovations, so that conflicts between goals are likely to arise. Apart from that, both characteristics may change in the course of time.

First of all, in this paper the characteristics of innovations, innovation processes, and innovation projects will be discussed in detail. Then, in order to deal with these particular characteristics, an approach will be taken which involves two planning instruments: On the one hand, a variant of stochastic network planning will be suggested, while on the other hand, this will be combined with multi-goal planning. This paper is aimed at presenting an instrument that supports the innovation planner interactively in handling challenging characteristics. In the discussion will be shown how aspects of uncertainty in connection with limited rationality occur in innovation management and what impact genetic algorithms may have.

Characteristics of Innovations, Innovation Processes, and Innovation Projects

To begin with, some characteristics of innovations, innovation processes, and innovation projects will be explained to provide a setting for the instruments discussed later on.

Innovations: In the parlance of business administration, an innovation is defined as a new product, a new service, a new production process, or a new combination thereof, which is introduced by a company for the sake of gaining a competitive advantage [MS05]. Accordingly, innovations comprise more than R&D activities, as marketing, production, purchasing, and other — especially managerial — activities have to be combined to create an innovation. As Kleinschmidt, Geschka and Cooper [KGC96] have shown by means of empirical analysis, the rate of success of innovation projects exceeds the average, if these activities are applied in parallel (with different intensity) throughout the entire innovation project. Two further characteristics are typical of innovations:

- (i) The innovator is faced with risks and uncertainties of different kinds and origins. Three major types of risks and uncertainties must be taken into account: technical risks, market risks, and resource risks [Spi06] (p. 31). For example: Will technology be ready and stable enough for the innovation? Will it be outpaced by another technology? Will the market demand still be there, when the development is complete? Or will it grow to new, unexpected

dimensions? Will the company be able to exploit a possible first mover advantage? Will competition between different investments within a company result in a resource stop for the innovation? — These risks and uncertainties have to be identified and considered during the planning process.

- (ii) The innovator has to deal with a multiplicity of goals, of which a few might complement each other, while many tend to be contradictory (for systems of goals see [SBA02] pp. 18-22, [MD05]). For instance, in the context of product innovation, lead time to market and the total cost of the innovation project are often negatively correlated. So are product quality or price per product unit. Some authors, like Brockhoff & Urban [BU88], try to subsume this multiple goal situation under the category of a net present value optimization, but for planning in accordance with flexibility issues it seems helpful to argue on the basis of disaggregated goals.

Innovation processes: An innovation process is defined as the generic structure of work packages and work flows within a company and beyond [DP95]. Such innovation processes can be defined on various levels, ranging from single business units to whole industries. For instance, due to legal requirements, there is a typical innovation process for companies in the pharmaceutical industry, as they are obliged to test substances in different phases in order to gain permission to market a new drug, c.f. [Ger04] pp. 24-25.

Innovation projects: Whereas an innovation process represents a generic structure, as defined above, an innovation project represents the particular structure of work packages and work flows for one specific innovation. Often, the innovation project is planned by means of adapting a generic structure, however, there still are many specific aspects to be taken into consideration. For instance, some innovation projects necessitate the involvement of suppliers, in which case an expansion of the generic structure may ensue.

Stochastic Networks for the Planning of Risks and Uncertainties

In the planning of innovation processes and projects, the first challenge is to survey and consider risks and uncertainties. An adequate instrument for this purpose is the technique of stochastic network planning, which will subsequently be presented in the form of its variant GERT (Graphical Evaluation and Review Technique), originated by Pritsker (1966), and the related software-tool GERT-Net [Sch02]. In GERT-aided process and project planning there are two distinct ways of approach to the examination of risks and uncertainties:

- One can either establish a stochastic branching with the logic of decision nodes
- or apply stochastic activity durations to the arcs.

The Logic of Decision Nodes and Arcs

In GERT-network graphs, decision nodes are employed to enable a branching of activities. The course of the process or project is defined by means of these decision nodes, each of which possesses an entrance logic and an exit logic. Three different elements are permissible as entrances, two as exits [Hen91] pp. 55-56, [Vö71] p. 26, and [Fix78] p. 16. Consequently, there are six different decision node types in total (see Fig. 3.50).

Three possible elements are at disposal on the entrance side:

- AND-entrance: The node is activated, as soon as all incoming activities are completed.

	exit	deterministic	stochastic
entrance			
AND			
inclusive-OR			
exclusive-OR			

Figure 3.50: Logic of Nodes. Source:[Hen91]

- Inclusive-OR-entrance: The node is activated, as soon as at least one incoming activity is completed.
- Exclusive-OR-entrance: The node is activated, as soon as exactly one incoming activity is completed.

On the exit side, there is a distinction of two elements:

- Deterministic exit: After the decision node has been activated, all outgoing activities are started, i.e. each activity possesses a probability parameter of 1.
- Stochastic exit: After the decision node has been activated, only one of the outgoing activities is performed. Which activity is to be performed, is decided on the basis of a pre-set discrete probability distribution.

The fact that the duration of activities can be a random variable, is another property of GERT-networks. This random variable is represented with the arcs and can be sustained by different types of probability distribution, e.g. exponential, binominal, normal or geometric distribution [WGW72] p.77. For instance, if an activity undergoes frequent repetition (e.g. in the course of a loop) one thus arrives at different values representing the actual duration of execution.

An Example from Automotive Industry

The following example of a complex innovation project from the automotive sector is meant to clarify how GERT modeling works. In the automotive industry great store is set by increasing the probability of a project's success as well as by reducing its duration. There are numerous arguments in favour of the curtailment of project duration [DKS97] p.98, e.g. regarding the decrease in the complexity of planning and controlling through the evasion of a temporal overlapping of projects, a reduction of the costs caused by capital employed, the attainment of new time margins and the minimization of difficulties with regard to schedule, and the establishment of opportunities and available capacities to facilitate a swifter execution of further projects in the future.

This particular aspect was treated by Schreyer [Sch99]: Starting from a GERT network graph, of the kind on which US-companies base their implementation of projects, Schreyer devised several variants, which differ from the basic network graph either due to the assumption of certain pre-conditions (reducing the duration of some activities by means of learning effects, innovation capacity and the involvement of repetitive elements) or to a consistent reformation of the process

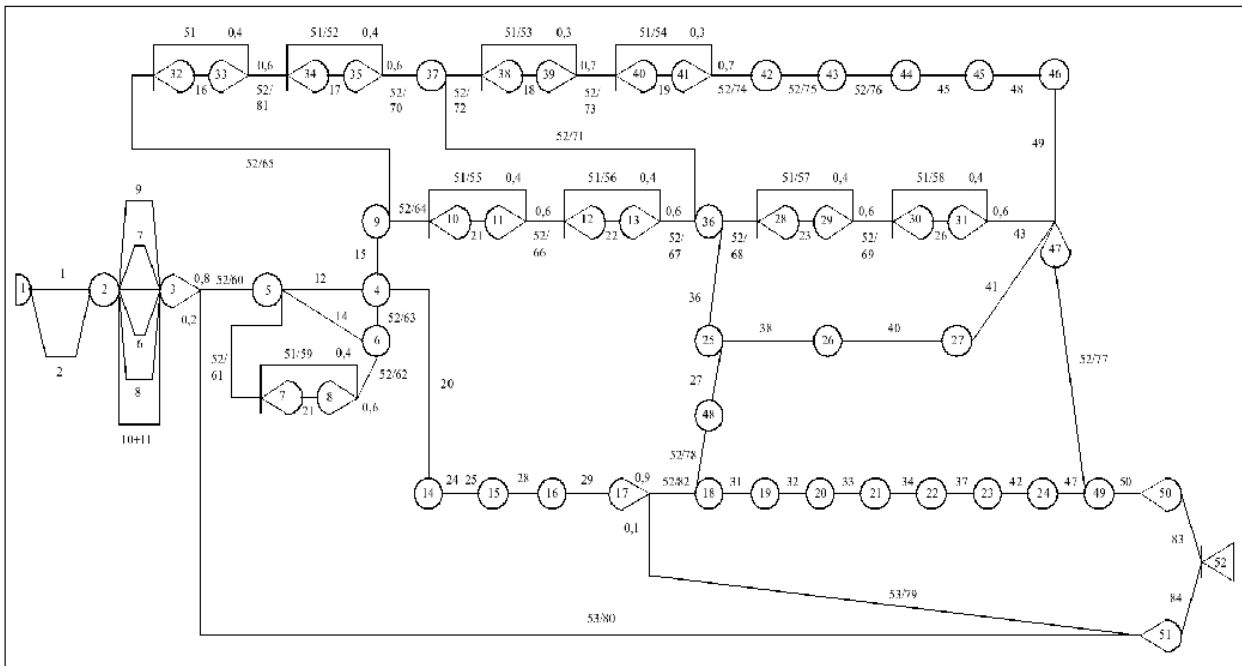


Figure 3.51: Stochastic network graph for an innovation in the automotive industry. For logic of nodes see figure 3.50. Source: [Sch99] p. 146.

structure (reengineering). One of these variants (see fig. 3.51) is characterized by a marked synchronization of processes [Sch99] p. 146.

The different variants were simulated with the software GERT-Net, which is used for the input, arrangement and simulation of GERT-network graphs [Sch02] (pp. 411-412). The calculation of the GERT-network graph for Schreyer’s reengineered variant produced a mean of 253 time units (4,8 years) for the entire project duration, with a probability of success amounting to 71 per cent (see fig. 3.52). The simulation of GERT-network variants with different parameterization, and a comparison of their results enable a detection of coherences and weaknesses within innovation processes projects.

Background Information on GERT

GERT-network planning represents a further development of well-known network planning techniques like CPM (Critical Path Method), PERT (Program Evaluation and Review Technique), MPM (Metra Potential Method) and GAN (Generalized Activity Networks), see fig. 3.53. The conventional, deterministic methods of network planning facilitate a delineation of projects, whose course is clearly determined as, during their realization, all activities are definitely completed [Fix78] p.9. Stochastic network graphs, however, are also suitable for surveying projects, whose course is not determined completely from the outset.

Planning with Multiple Goals

In addition to the consideration of different risks and uncertainties, the pursuance and coordination of multiple goals represent another challenge in the planning of innovation processes and projects. Subsequently, different systems of goals will be introduced, followed by a discussion of Pareto-optimization within such systems.




Figure 3.52: Frequency distribution of time required for the final node in case of project success. Source: [Sch99] p. 157.

Duration of activities \ Sequence of activities	deterministic	stochastic
	deterministic	CPM, MPM
stochastic	PERT	GERT

Figure 3.53: Overview of deterministic and stochastic network planning techniques. Source: [Sch99] p. 51

Systems of Goals for Innovation Projects

In business management, goals are synonymous with desired future conditions in or of an organization. As the definition thereof often relates to various components, and as decision-makers tend to differ in their understanding of the assumed goals, these future conditions are described as systems of goals or multiple goal decisions [BS02a]. In innovation projects, the goal system is often defined with reference to the project management system of goals, which contains the so-called „magic triangle“ of quality, cost and schedule [Har04]. This may be useful, as it involves the application of proven systematics. However, it also stands for a limitation of creative scope within an innovation project, as different goals either have to be converted into restrictions or be included in the quality target by way of a benefit calculation.

Therefore, it seems appropriate to plead for the use of a more comprehensive system of goals, such as the one devised by Specht, Beckmann and Amelingmeyer [SBA02] p. 19. Here, each goal included in the system contributes to the attainment of economic and non-economic targets within the company. The goals that are relevant for an innovation project can, for example, be sorted into three groups: project-related, product-related, and production-related goals (see fig. 3.54).

The process of defining goals is a vital component of innovation planning. Based on the company’s aims, innovation planning has to include the determination and formulation of specific goals in accordance with the given situation. Consequently, planners are required to monitor their system of goals constantly, in order to maintain its adaptability throughout the course of the project [SBA02] pp. 19 and 22.

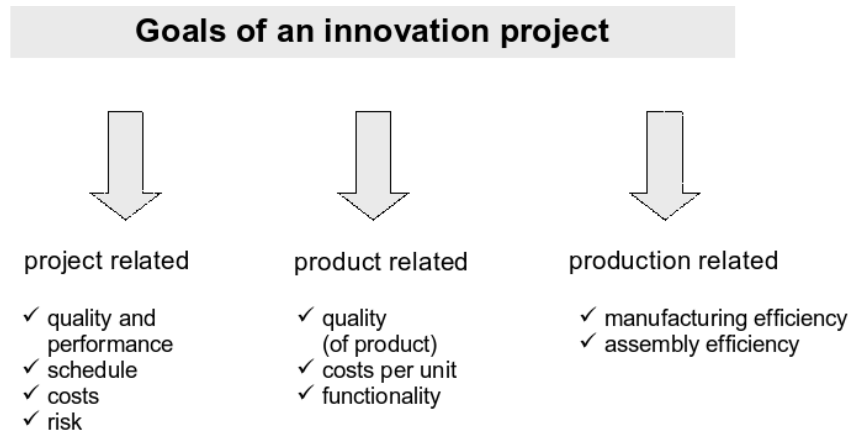


Figure 3.54: Concrete goal system of an innovation project. Source: [MD05] p. 231

The outlined system of goals for innovation projects does not result in all goals being complementary (for a similar product life cycle study see [Dow01]). In fact, there are several contradictions. Some examples:

- In most cases the costs of an innovation project, which ends with the market launch of a new product, and the costs per product unit, will evolve in opposite directions: Extensive development and testing are likely to generate lower costs per product unit, but higher project costs.
- The schedule of an innovation project and the functionality of its new product will also be in opposition: The sooner an innovation project has to be completed, the fewer functions its product is likely to possess.
- The same relationship can be found between the schedule of an innovation product and the manufacturing efficiency of the related product.

Pareto Optimization

Planning with multiple goals means seeking solutions to problems in which several competing goals are required to be optimized, see [Gö95] p. 96, and the articles of Küfer (pages 77 ff.) and Roosen (pages 111 ff.) in this book. Generally, this allows for no such thing as a „perfect“ solution, which — in the case of simple problems — is determined by means of a combination of alternatives, and in which all goal functions assume their extreme values [Vie82] p. 194. Instead, one tries to find a solution set with regard to different functions of goals, i.e. various solutions that are, in a sense, just as valuable as others. Accordingly, there is no „best“ solution to speak of — only one variant over which others have no predominance.

This type of solution is generated with the aid of Pareto optimization [Gö95] p. 96. The concept of Pareto dominance was devised by Vilfredo Pareto, who laid the foundations of research into multi-criteria optimization. A target vector with certain values dominates another, if it reaches a higher value for at least one criterium while assuming no lower value for any of the remaining criteria. Consequently, a Pareto optimal set consists of solutions which are not dominated by others. The algorithms used in Pareto-based techniques are aimed at calculating and offering a set of excellent and, preferably, diverse solutions. Then, users are able to select a solution from this set according to their individual preferences.

For instance, the management of a company who have commissioned an innovation project can discuss options of handling the conflicting goals of project costs and costs per product unit with the project manager (see fig. 3.55). This produces a Pareto set, which does not contain a single „best“ point of realization, but requires the management to make a decision in favour of one predominant point on the basis of additional considerations.




Figure 3.55: Pareto-set for the comparison of project costs and costs per product unit

A Concept of Integration: The Stochastic, Multi-Goal-Oriented, Interactive Planning of Innovation Projects.

With stochastic network planning on the one hand, and multi-goal-planning on the other, there are two instruments at disposal, which each permit a separate handling of the aforementioned challenges to innovation planning. What is still missing, is an integration of these instruments in connection with the consideration of changes in the schedule. To deal with this, we suggest a stochastic, multi-goal-oriented and interactive planning process. Here, three theoretical foundations have to be unified. (i) Stochastic network graphs with expanded parameterization are required for the modelling of innovation projects. (ii) Genetic algorithms help to establish Pareto-sets and thus assist a multi-goal optimization. (iii) Decision-makers must be enabled to influence the course of the process or project interactively, especially concerning the definition of a system of goals.

One possible procedure for the optimization of innovation projects will subsequently be presented (see 3.56). It may be divided into three phases: (i) a basic simulation, (ii) a Pareto optimizing simulation, and (iii) interactive changes in the settings after changes in the environment.

Basic simulation: For the basic simulation the planning problem is first modelled by means of a GERT-network graph with decision nodes and stochastic arcs. If a generic structure for an innovation process exists, it may be used as foundation. Relationships between activity duration and goals have to be specified (see below), and activity acceleration figures are set to 0 (see also below). Then the simulation model is repeatedly run. With GERT simulations it can be calculated

- how probabilities are distributed for success and abandonment of the project respectively,
- how project duration and other variables, which depend on the number of iterations within loops, are distributed in the cases of success and abandonment respectively,

- and how the remaining variables for the goals are in the basic attempt (in those cases one will find some upper or lower bounds e.g. for costs per product unit).

The basic simulation helps understanding the innovation project and evaluating the parameters used.

Pareto optimizing simulation: Then, the innovation project is to be optimized with regard to its defined goals. This includes a variation of activity acceleration figures and the establishment of relative degrees of goal realization. With the aid of genetic algorithms, Pareto sets of the respective results are drawn up and presented to the user in a comprehensible graphic form. From these Pareto sets the user selects the solution which is best suited to the current situation.

Interactive changes: In the course of an innovation project, goals may change dramatically because of competitor activities, the restructuring of resources or political circumstances. In such cases, the innovation project planning has to be updated by means of appropriate re-optimization.

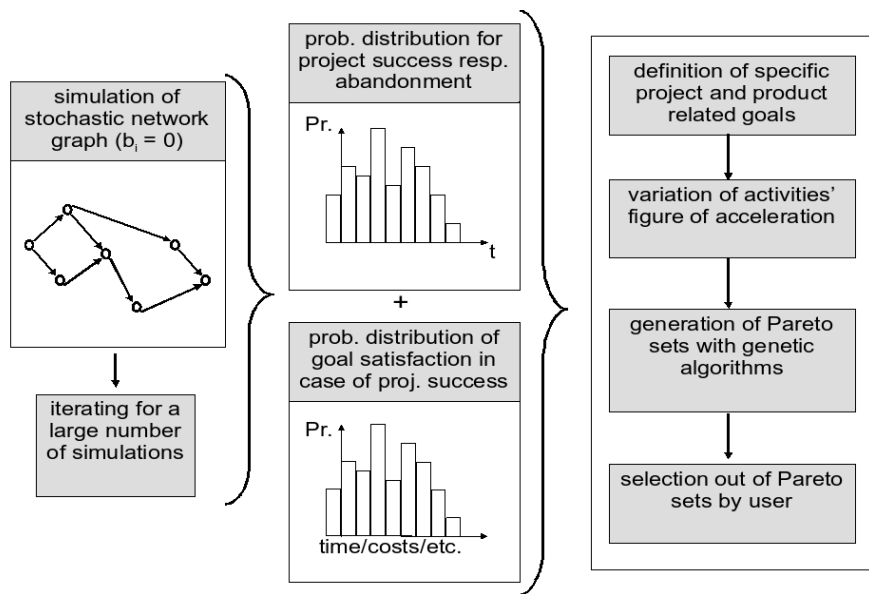


Figure 3.56: Possible procedures for the optimization of innovations

In the following subsections, two idiosyncrasies of the model of procedure illustrated above will be dealt with: First, we will attend to GERT-network planning with expanded parameterization. Second, we will take a closer look at how adequate solutions are generated and selected, and discuss possible ways of representing Pareto solutions.

Expanded Parameterization of GERT-networks with Goal Functions

In accordance with the suggested procedure, a module should be created on the basis of the GERT-Net program, which permits a comprehensive parameterization of activities in an innovation project. In addition to the usual GERT-network activity parameters, i.e. the probability of the completion of an activity p_i and the duration of the activity d_i , two further parameters are to be defined for this purpose: (i) the specific acceleration figure of the activity b_i and (ii) its specific relation with goals $f(a_i)$ (see fig. 3.57).

Acceleration figure: The management is able to influence the duration of individual activities decisively by various measures, which might cause a change in costs. Their scope for action includes the selection of a specific acceleration figure, which shows to what percentage the temporal acceleration potential is being exploited [Sac00] pp. 135-136. The duration of an activity

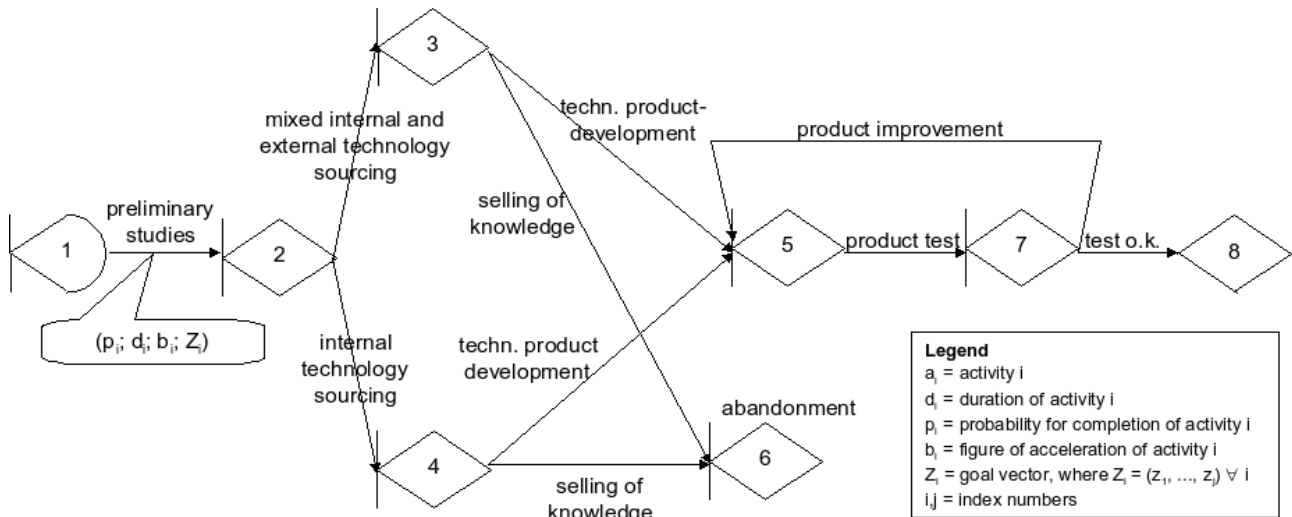


Figure 3.57: Example of GERT-network graph with expanded parameterization

can be reduced, e.g. by means of new methods, innovation improvement based on learning effects, the improvement of internal organizational coordination and communication, the utilization of additional resources or outsourcing [Sch99] p. 133.

Relations with goals: An activity's specific relation with goals is illustrated by means of functions, and shows to what extent each predetermined goal is influenced by the related activity. To give an example, a possible connection between the duration of product development and the costs per product unit is illustrated below (see fig. 3.58). This relation is partly marked by reverse trends. In the first section, there is a clear opposition: The more intensive and, consequently, extensive the phase of development is, the lower the costs per unit are likely to be. In a second section, which lies beyond a certain quantity of time for product development, no further effectuation can be detected.

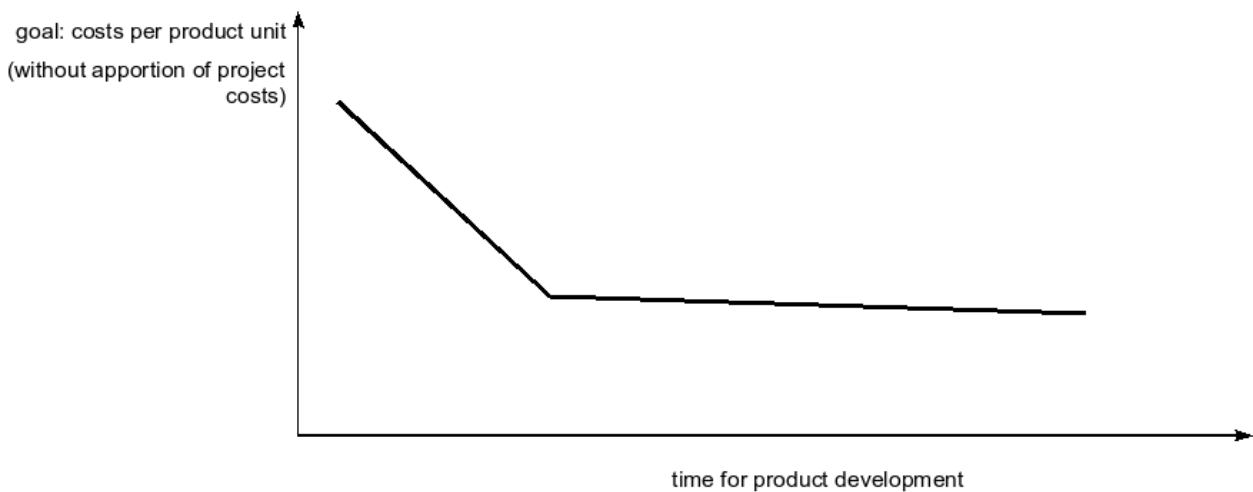


Figure 3.58: Specific relation between duration of product development and costs per product unit

Generation and Selection of Pareto optimal Solutions

Algorithms that employ Pareto based techniques are aimed at calculating and offering a set of solutions which are not only excellent but also as diverse as possible. Those solutions which are not dominated by others, are then presented to the decision-maker. For this purpose, the suggested procedure should be incorporated in a graphic user interface for the figuration of Pareto sets. Such a clear pictorial display is supposed to assist users in making their decisions.

In addition to project related goals, project costs and schedule, for instance, product related goals such as quality, safety and service options are to be taken into account. Some of these goals can be predetermined: In spite of minimal project costs a high quality and a first-rate service are to be attained. All remaining goals are subject to variation. This procedure is illustrated in fig. 3.59.

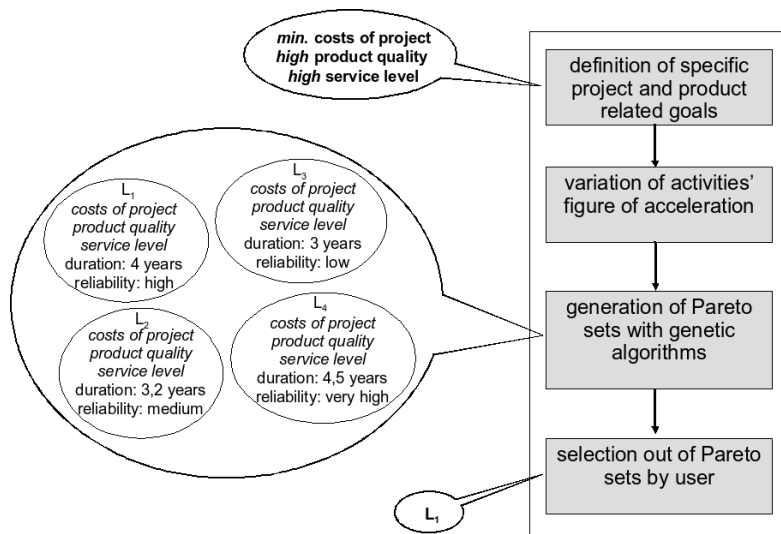


Figure 3.59: Generation and selection of Pareto sets. The predefined goals within the Pareto sets are printed in italics.

The generated Pareto sets are now to be graphically represented. In this particular case a two-dimensional representation will suffice, as three of five goals are predetermined. If a great number of goals and, consequently, a great number of dimensions are involved, it becomes particularly important to provide a clear graphic display. With more than three goals subject to variation at the same time, an interactive display has to be employed [dFIfTuW03].

Discussion

To facilitate a discussion, two aspects treated by the BBAW study group will be singled out: (i) uncertainty in connection with limited rationality, and (ii) the impact of genetic algorithms.

- As the previous chapters have shown, uncertainties and risks are of crucial importance where innovations are concerned. This fact leads to the question, how intensively innovation planners are using stochastic network planning at present. The reply is that they still retain an attitude of reserve towards this method of modeling, not entirely without reason. One cause may be that, so far, no adequate professional software tool is available. Moreover, some sort of limited rationality may be at play: Possibly, innovation planners tend to avoid stochastic network planning for reasons of communication. To be more precise: If an innovation planner has to settle a date for a specific task with an employee, it appears easier to agree on a certain day than on a period of time. Hence, innovation planners might still

prefer the less complex method of deterministic network planning (see chapter 3.8), although stochastic network planning would provide far better insights into the time and risk frame of the innovation project. In order to overcome this planning dilemma, a combination of both variants appears to be sensible.

- Evolutionary strategies for optimization may be of great help in solving problems with innovation planning. They are especially valuable: (i) if there are great number of non-linear relations between activities and goals, and (ii) if the goals of an innovation project are shifting — which may be due to various causes, such as competitors' activities or legal changes — a rapid re-optimization of the innovation project has to take place. In such cases, the specific attributes of genetic algorithms [Rec94] p. 218 can be helpful.

In general, the presented concept is an approach towards a new understanding of the structures of innovation processes and projects. However, it requires further development. Also, prototypes of a respective IT system should be implemented and tested.

3.9 Bounded Rationality: Fast and Frugal Heuristics for Decisions under Uncertainty

Konstantinos V. Katsikopoulos and Gerd Gigerenzer

To the great relief of passengers all around the world airplanes take off, fly, and land every day. Mathematicians, physicists, mechanical and aeronautical engineers have created a large store of know-how on the technical aspects of aviation. The management of the aviation industry, however, remains challenging. Consider, for example, making decisions under uncertainty. We do not know how to route passengers to jointly optimize time, safety, and fuel efficiency. And, even if we focused on just one objective, say time, the routing problem is computationally intractable.

Like in aviation, many decision problems in large-scale engineering systems resist formulation in optimization terms because they have multiple criteria and require information that is not easily available (e.g., utilities and probabilities). Even when they can be formulated as optimization problems, computing an optimal solution often turns out to be intractable. What to do? The approach typically taken in the natural sciences and engineering is to optimize a simplification of the original problem.

This often works well but it can be difficult to know if the simplifications hold and what is the loss when they are violated. More broadly, the optimization approach — with some exceptions, e.g., Taguchi methods — is not tailored to handle issues like uncertainty, robustness, and flexibility that are increasingly recognized as fundamental in the management of engineering systems [Pap04]. Finally, the interface between optimization methods and practitioners often lacks transparency, usability, and acceptance [KOCZ93].

Overall, engineers are trained in the rigorous theory of optimization but all too often, when they graduate, they seem to find themselves making do by using heuristics that worked in the past. At the Max Planck Institute for Human Development in Berlin, our team of life, natural, and social scientists, as well as historians, philosophers, mathematicians, and engineers, has, for more than ten years, been studying heuristics for decision making under uncertainty. Our research program can be viewed as the study of bounded rationality, a popular concept in the social sciences [Sim55]. One of our interests is in modeling the heuristic cognitive processes laypeople and some practitioners — medical doctors and mock jurors — use. This chapter samples some answers to the normative question of how well do the heuristics perform. The research has used decision-making tasks that do not directly relate to engineering concerns but we also speculate how it can be applied to the management of engineering systems. We start with some general comments on the heuristic view of bounded rationality.

Bounded Rationality: Fast and Frugal Heuristics

Bounded rationality is often interpreted as optimization under constraints, where the constraints are due to impoverished cognitive ability or incomplete information [Con96]. This interpretation is contrary to Simon's (1955) who emphasized satisficing, or picking any outcome that exceeds a pre-determined aspiration level (as opposed to picking only an optimal outcome). Furthermore, this interpretation does not allow much progress as it reverts the study of decision making back to the study of logic, probability, and calculus, while excluding psychology.

Our approach to bounded rationality takes an ecological rather than logical view. It does not study optimal, internally consistent decisions but decisions that surpass aspiration levels with regard to external criteria like speed, accuracy, robustness, and transparency. This fits well with

engineering where the focus is not so much on internal consistency but on external performance. Gigerenzer, Todd, and colleagues [GTtARG99, TGtArgip19] model the decisions with a breed of simple rules of thumb, called *fast and frugal heuristics*, which use a minimum of time, information, and computational resources.

The heuristics can be understood from a Darwinian perspective. First, because evolution does not follow a grand plan, there is a patchwork of heuristics, tailored to particular problems. This gives flexibility to the bounded rationality approach. Second, just as evolution produces adaptations that are bound to a particular ecological niche, the heuristics are not rational or irrational, per se, but only relative to an environment. Note that the study of the interaction between the decision-maker and the environment — emphasized in the ecological approach to psychology [Bru55, Sim56] — is missing in the optimization-under-constraints approach.

Finally, and importantly, heuristics exploit core psychological capacities, like the ability to track objects. This is exactly what allows the heuristics to be simple, yet successful. For example, consider a pilot who spots another plane approaching, and fears a collision. How can she avoid it? A simple heuristic that works is to look at a scratch in her windshield and observe whether the other plane moves relative to that scratch. If it does not, she should dive away quickly.

In short, in our view, bounded rationality is about simple and transparent heuristics that require minimum input, do not strive to find a best and general solution, but nevertheless are accurate and robust. What do these heuristics look like?

The Recognition Heuristic

Imagine you are a contestant in a TV game show and face the \$1 million question: Which city has more inhabitants: San Diego or San Antonio?

What is your answer? If you are American, then your chances of finding the right answer, San Diego, are not bad. Some two thirds of undergraduates at the University of Chicago did [GG02]. If, however, you are German, your prospects look dismal because most Germans know little about San Diego, and many have not even heard of San Antonio. How many correct inferences did the less knowledgeable German group that we tested achieve? Despite a considerable lack of knowledge, 100% of the Germans answered the question correctly. How can people who know less about a subject nevertheless make more correct inferences? The answer seems to be that the Germans used a heuristic: If you recognize the name of one city but not the other, then infer that the recognized city has the larger population. The Americans could not use the heuristic, because they had heard of both cities. They knew too much.

The recognition heuristic is useful when there is a strong correlation — in either direction — between recognition and criterion. For simplicity, we assume that the correlation is positive. For paired-comparison tasks, where the goal is to infer which one of two objects (e.g., cities) has the higher value on a numerical criterion (e.g., population), the following fast and frugal heuristic can be used.

Recognition heuristic: If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value on the criterion.

The recognition heuristic builds on the core capacity of recognition — such as face, voice, and name recognition. No computer program yet exists that can perform face recognition as well as a human child does. Note that the capacity for recognition is different from that for recall. For instance, one may recognize a face but not recall anything about who that person is [CM87].

Intuitively, the recognition heuristic is successful when ignorance is systematic rather than random, that is, when recognition is strongly correlated with the criterion. The direction of the correlation between recognition and the criterion can be learned from experience, or it can be genetically coded. Substantial correlations exist in competitive situations, such as between name

recognition and the excellence of colleges, the value of the products of companies, and the quality of sports teams.

Consider forecasting the outcomes of the 32 English F.A. Cup third-round soccer matches, such as Manchester United versus Shrewsbury Town. Ayton and Önkal (1997) tested 50 Turkish students and 54 British students. The Turkish participants had very little knowledge about (or interest in) English soccer teams, while the British participants knew quite a bit. Nevertheless, the Turkish forecasters were nearly as accurate as the English ones (63% versus 66% correct). Their predictions were consistent with the recognition heuristic in 627 out of 662 cases (95%). More generally, a number of experimental studies have found that if the accuracy of the recognition heuristic, α , is substantial (i.e., exceeds, say, 0.7), people use the recognition heuristics in about 90% of all cases [Nay01].

The recognition heuristic implies several counterintuitive phenomena of human decision making that cannot be deduced from any other theory we are aware of. For instance, recognition information tends to dominate further knowledge, in rats as well as in people, even if there is conflicting evidence [GG02]. Here we concentrate on the counterintuitive finding that less information can increase accuracy.

The Less-is-More Effect

Assume that a person recognizes n out of N objects. The probability of being able to use the heuristic equals the probability of exactly one object being recognized, or

$$r(n) = \frac{2n(N-n)}{N(N-1)}. \quad (3.9)$$

Similarly, the probability that both objects are recognized, and thus other knowledge beyond recognition must be used equals

$$k(n) = \frac{n(n-1)}{N(N-1)}. \quad (3.10)$$

Finally, the probability that neither object is recognized, which is taken to mean that the person has to guess, equals

$$g(n) = \frac{(N-n)(N-n-1)}{N(N-1)}. \quad (3.11)$$

Let β be the accuracy of the person when both objects are recognized and other knowledge is used. We also assume that accuracy equals $\frac{1}{2}$ when none of the objects is recognized (and we also assume $\alpha, \beta > \frac{1}{2}$). Thus, the overall accuracy of a person who recognizes n objects equals

$$f(n) = r(n)\alpha + k(n)\beta + g(n)\left(\frac{1}{2}\right). \quad (3.12)$$

Definition 1. The *less-is-more effect* occurs when there exist n_1 and n_2 so that $n_1 < n_2$ but $f(n_1) > f(n_2)$ with $n_1, n_2 \in 0, 1, \dots, N$.

Definition 2. The *prevalence*, p , of the less-is-more effect is the proportion of pairs (n_1, n_2) with $n_1 \neq n_2$ for which the less-is-more effect occurs.

The prevalence p of the less-is-more effect varies between zero — no effect — for increasing $f(n)$, and unity — there is always an effect — for strictly decreasing $f(n)$. The prevalence of the less-is-more effect depends on the person's α and β . For $\alpha = 0.8$ and $\beta = 0.6$, simple enumeration yields $p = 1/3$. More generally, the following holds [RK04].

Result 1. The less-is-more effect occurs (i.e., $p \neq 0$) if and only if $\alpha > \beta$. The effect becomes more prevalent (i.e., p increases) as α increases or β decreases.

At a first glance, the less-is-more effect might appear paradoxical. But it is not, because less recognition information may simply enable more accurate cognitive processing (via the use of the recognition heuristic). This is formalized by $\alpha > \beta$.

As an example, Goldstein and Gigerenzer discuss three Parisian sisters who have to compare all pairs of cities from the most populous $N = 100$ German cities [GG02]. All sisters have $\alpha = 0.8$ and $\beta = 0.6$, but they vary on the number of recognized objects: The youngest sister has $n = 0$, the middle sister has $n = 50$, and the eldest sister has $n = 100$. When $\alpha > \beta$ the strong less-is-more effect is predicted: for the middle sister, $f(50) = 0.68$, while for the eldest sister $f(100) = 0.60$. Accuracy for $\alpha = 0.8$ and $\beta = 0.6$, based on eqs. 3.9 – 3.12, interpolated for all n , is graphed in fig. 3.60 (solid curve; the dashed curve will be explained below).

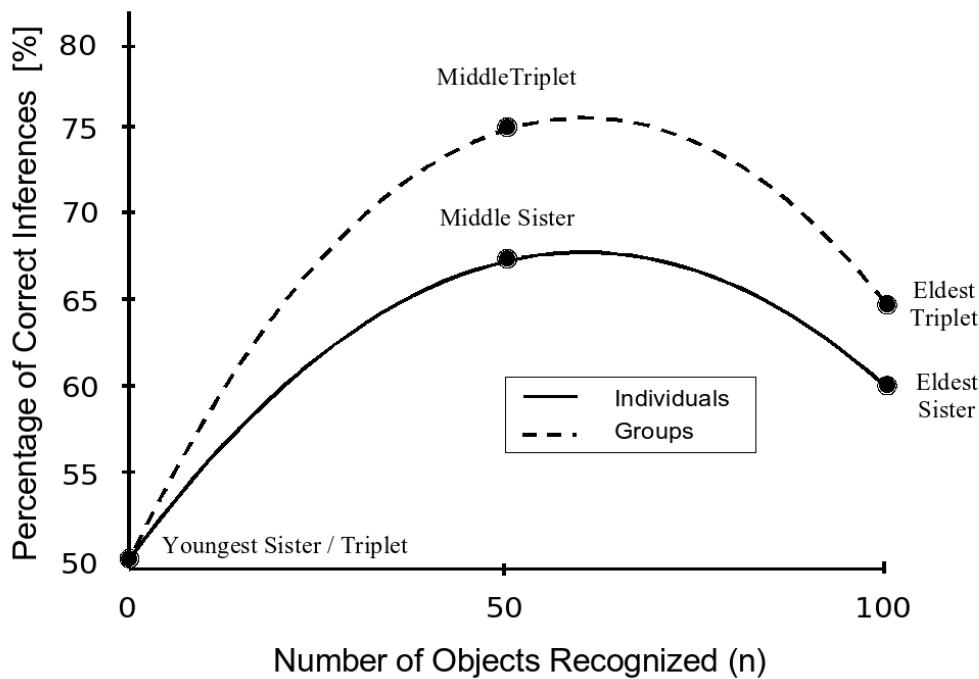


Figure 3.60: Predicted accuracy as a function of number of recognized objects for $\alpha = 0.8$ and $\beta = 0.6$, for individuals (solid curve) and three-member groups that use the majority rule (dashed curve).

A less-is-more effect can emerge in at least three different situations. First, it can occur between domains, that is, when the same group of people achieves higher accuracy in a domain in which they know little than in a domain in which they know a lot. For instance, when American students were tested on the 22 largest American cities (such as New York versus Chicago) and on the 22 most populous German cities (such as Cologne versus Frankfurt), they scored a median 71.0% (mean 71.1%) correct on their own cities but slightly higher on the less familiar German cities, with a median of 73.0% correct (mean 71.4%). This effect was obtained despite a handicap: Many Americans already knew the three largest U. S. cities in order, and did not have to make any inferences [GG02]. A similar less-is-more effect was demonstrated with Austrian students, whose scores for correct answers were slightly higher for the 75 largest American cities than for the 75 largest German cities [Hof95]. Second, a less-is-more effect can occur during knowledge acquisition, that is, when an individual's performance curve first increases but then decreases again. Finally, the effect can occur between two groups of people, when a more knowledgeable group

makes worse inferences than a less knowledgeable group in a given domain. An example is the performance of the American and German students on the question of whether San Diego or San Antonio is more populous [GG02]. Furthermore, Reimer and Katsikopoulos [RK04] ran a study where groups of people decided together.

In this study, three people sat in front of a computer screen on which questions such as “Which city has more inhabitants: San Diego or San Antonio?” were displayed. The task of the group was to find the correct answer through discussion, and they were free to use whatever means. The correct solution is difficult to prove by an individual group member, thus one might expect that the majority determines the group decision [GH97]. The accuracy, $G(n)$, of a group using the majority rule is calculated as follows.

Assume first that the group is *homogeneous* (i.e., all members have equal α , β , and n) and *independent* (i.e., the recognition and inference processes of members are independent given the values of the criterion on the objects). Let $F(i)$ be the probability of exactly i members being accurate and the group, using the majority rule, being correct. Finally let the group have m members, $c(m, i)$ be the number of ways in which i objects can be sampled out of m objects without replacement, and $\text{majority}(m) = (m + 1)/2$ if m is odd, and $= (m/2 + 1)$ if m is even. Then, the following holds:

$$\begin{aligned} F(i) &= c(m, i) f(n)^i (1 - f(n))^{m-i} (1/2), \quad i = \text{majority}(m) \quad \& \quad m \text{ is even} \\ &= c(m, i) f(n)^i (1 - f(n))^{m-i}, \quad \text{otherwise.} \end{aligned} \quad (3.13)$$

$$G(n) = \sum_{i=\text{majority}(m), \dots, m} F(i). \quad (3.14)$$

The application of (3.13) and (3.14) for $\alpha = 0.8$, $\beta = 0.6$, and $m = 3$ is illustrated in Figure 1 (dashed curve). A less-is-more effect is again predicted and $p = 1/3$. More generally, the following holds [RK04].

Result 2. Assume a homogeneous and independent group, using the majority rule. Then, (i) less-is-more effect is predicted if and only if $\alpha > \beta$ and (ii) the prevalence of the effect equals the prevalence for one member.

Consider now the following conflict. Two group members have heard of both cities and each concluded independently that city A is larger. But the third group member has not heard of A, only of B, and concludes that B is larger (relying on the recognition heuristic). After the three members finished their negotiation, what will their consensus be? Given that two members have at least some knowledge about both cities, one might expect that the consensus is always A, which is also what the majority rule predicts. In fact, in more than half of all cases (59 %), the group voted for B [RK04]. This rose to 76 % if two members used recognition.

Group members letting their knowledge be dominated by others' lack of recognition may seem odd. But in fact this apparently irrational decision increased the overall accuracy of the group. Broadly consistent with Result 2, Reimer and Katsikopoulos [RK04] observed that, when two groups had the same average α and β (that were such that $\alpha > \beta$), the group who recognized fewer cities (smaller n) typically had more correct answers. For instance, the members of one group recognized on average only 60 % of the cities and those in a second group 80 %, but the first group got 83 % answers correct in a series of over 100 questions, whereas the second only 75 %. Thus, group members seem to intuitively trust the recognition heuristic, which can improve accuracy and lead to the counterintuitive less-is-more effect between groups.

Cue-Based Heuristics

When recognition is not valid, or people recognize all objects, heuristics can involve search for reasons or, in psychological jargon, *cues*. A few years after his voyage on the *Beagle*, the 29-year-old Charles Darwin divided a scrap of paper (titled, “This is the Question”) into two columns with the headings “Marry” and “Not Marry” and listed supporting reasons for each of the two possible courses of action, such as “nice soft wife on a sofa with good fire” opposed to “conversation of clever men at clubs.” Darwin concluded that he should marry, writing “Marry – Marry – Marry Q. E. D” decisively beneath the first column [Dar69] p. 232–233. The following year, Darwin married his cousin, Emma Wedgwood, with whom he eventually had 10 children. How did Darwin decide to marry, based on the possible consequences he envisioned — children, loss of time, a constant companion? He did not tell us. But we can use his “Question” as a thought experiment to illustrate various visions of decision making.

Darwin searched in his memory for reasons. There are two visions of search: optimizing search and heuristic search. Following Wald’s (1950) optimizing models of sequential analysis, several psychological theories postulated versions of sequential search and stopping rules [BT93]. In the case of a binary hypothesis (such as to marry or not marry), the basic idea of most sequential models is the following: A threshold is calculated for accepting one of the two hypotheses, based on the costs of the two possible errors, such as wrongly deciding that to marry is the better option. Each reason or observation is then weighted and the evidence is accumulated until the threshold for one hypothesis is met, at which point search is stopped, and the hypothesis is accepted.

If Darwin had followed this procedure, he would have had to estimate, consciously or unconsciously, how many conversations with clever friends are equivalent to having one child, and how many hours in a smoky abode can be traded against a lifetime of soft moments on the sofa. Weighting and adding is a mathematically convenient assumption, but it assumes that there is a common currency for all beliefs and desires in terms of quantitative probabilities and utilities. These models are often presented as models whose task is to predict the outcome rather than the process of decision making, although it has been suggested that the calculations might be performed unconsciously using the common currency of neural activation.

The second vision of search is that people use heuristics — either social heuristics or cue-based heuristics — that exploit some core capacities. Social heuristics exploit the capacity of humans for social learning and imitation (imitation need not result in learning), which is unmatched among the animal species. For instance, the following heuristic generates social facilitation [Lal01]:

Do-what-the-majority-do heuristic: If you see the majority of your peers display a behavior, engage in the same behavior.

For the marriage problem, this heuristic makes a man start thinking of marriage at a time when most other men in one’s social group do, say, around age 30. It is a most frugal heuristic, for one does not even have to think of pros and cons. Do-what-the-majority-do tends to perform well when (i) the observer and the demonstrators of the behavior are exposed to similar environments that (ii) are stable rather than changing, and (iii) noisy, that is, where it is hard to see what the immediate consequence of one’s action is [BR85] [GGH⁺01].

Darwin, however, seems to have based his decision on cues. We will describe two classes of heuristics that search for cues. Unlike optimizing models, they do not weight and add cues. One class of heuristics dispenses with adding, and searches cues in order (a simple form of weighing). These are called *lexicographic heuristics*. Another class dispenses with weighting and simply adds, or *tallies*, cues.

Lexicographic and Tallying Heuristics

We again consider the comparison task in which both objects are recognized. The decision is made on the basis of n binary cues c_1, c_2, \dots, c_n , $n \geq 2$ (for any object a cue c_i equals 1 or 0). Since the 18th century, a popular decision rule for paired comparisons is the linear rule [KMM02] (2003). In this rule, cue c_i has a weight w_i ; weights can be estimated by, say, minimizing the sum of squared differences between the predictions of the rule and the observations. For an object A with cue values $c_i(A)$, the score $\sum_i c_i(A)w_i$ is computed and the object with the higher score is picked. If the scores are equal, an object is picked randomly. Tallying is a linear rule where weights are equal to unity, an old idea in psychological measurement [Gul50].

Take the best is a heuristic in the lexicographic tradition, which dates back to at least thirty-five years [Tve72] (1969). First, cues are ordered by decreasing validity, where the validity v_i of cue c_i is the conditional probability that the cue points to the larger object ($c_i = 1$ on the larger object and $c_i = 0$ on the other object) given that the cue discriminates between the objects [$c_i(A) \neq c_i(B)$]. (Without loss of generality it can be assumed that $1 \geq v_i \geq \frac{1}{2}$). After cues are ordered, the decision maker inspects the first cue. If this cue points to one of the objects then this object is taken to be larger. If the cue does not discriminate between the objects, then the second cue is inspected and so on until a discriminating cue is found; if no such cue exists, an object is picked at random.

One-cue decision making has been observed in high-stake decisions. British magistrates tend to make bail decisions on the basis of one good reason only [Dha03, DA01], and so do British general practitioners when they prescribe lipid-lowering drugs [DH01]. Many parents rely on one cue to decide which doctor to drive to in the night when their child becomes seriously ill [Sco02].

Both *take the best* and *tallying* are naïve in the sense that they do not take cue dependencies into account. While at a first glance they might appear simplistic, simulation studies have shown that naïve heuristics compare remarkably well to statistical benchmarks. Three decades ago, Dawes and Corrigan (1974) convincingly argued that tallying can have greater predictive accuracy than linear regression. Einhorn and Hogarth (1975) provided statistical reasons for this, like the absence of sampling error in the estimation of weights. Czerlinski, Gigerenzer, and Goldstein (1999) replicated this finding in twenty real-world datasets, emphasizing the concepts of *overfitting* and *robustness*. To define these, we distinguish between a learning sample from which a model estimated its parameters and the test sample on which it is tested. Both samples are randomly drawn from the same population.

Definition 3. A model M over-fits the learning sample if there exists an alternative model M' such that M has a smaller error than M' in the learning sample but a larger error in the test sample. In this case, M' is called the more robust model.

Figure 3.61 shows the accuracy of three heuristics compared to linear regression, averaged across 20 real-world problems [CGG99], e.g., to predict which Chicago public high school has the higher dropout rate based on the socioeconomic and ethnic compositions of the student bodies, the sizes of the classes, and the scores of the students on various standardized tests. (Other problems were to predict people's attractiveness judgments, homelessness rates, adolescents' obesity at age 18, etc.) The three heuristics were *take the best*, *minimalist* (which is a lexicographic heuristic that searches cues in random order), and *tallying*. *Take the best* and *minimalist* were most frugal; they looked up, on average, only 2.4 and 2.2 cues before they stopped search. *Tallying* and *multiple regression* looked up all cue information, which amounted to an average of 7.7 cues. How accurate are the heuristics?

Linear regression had the best fit. However, the true test of a method concerns its predictive accuracy, which was tested by cross-validation, that is, the four methods learned their parameters on half of the data (learning sample), and were tested on the other half (test sample). Figure 2 shows that *regression* over-fitted the data relative to both *take the best* and *tallying*. An intuitive




Figure 3.61: Robustness of three heuristics and linear regression, averaged across 20 real-world problems. Confidence intervals were, for all problems and methods, negligible

way to understand overfitting is the following: A set of observations consists of information that generalizes to the other samples, and of information that does not (e.g., noise). If one extracts too much information from the data, one will get a better fit but one will mistake more noise for predictive information. The result can be a substantial decrease in one's predictive power. Note that both forms of simplifying — dispensing either with adding or with weighting — resulted in greater robustness. *Minimalist*, however, which dispenses with both weighting and adding, extracts too little information from the data.

In general, the predictive accuracy of a model increases with its fit, decreases with its number of adjustable parameters, and the difference between fit and predictive accuracy gets smaller with larger number of data points [Aka73, FS94]. The general lesson is that in judgments under uncertainty, one has to ignore information in order to make good predictions. The art is to ignore the right kind. Heuristics that promote simplicity, such as using the best cue that allows one to make a decision and ignore the rest, have a good chance of focusing on the information that generalizes.

These results may appear counterintuitive. More information is always better; more choice is always better — so the story goes. This cultural bias makes contrary findings look like weird oddities [HT04]. Yet experts base their judgments on surprisingly few pieces of information [Sha92], and professional handball players make better decisions when they have less time [JR03]. People can form reliable impressions of strangers from video clips lasting half a minute [AR93], shoppers buy more when there are fewer varieties [IL00], and zero-intelligence traders make as much profit as intelligent people do in experimental markets [GS93]. Last but not least, satisficers are reported to be more optimistic and have higher self-esteem and life satisfaction, whereas maximizers excel in depression, perfectionism, regret, and self-blame [SWM⁺02]. Less can be more.

Beyond computer simulations, mathematical analyses have also been used to investigate the accuracy of heuristics [MH02, HK05, KM00]. In the case where cues are conditionally independent (i.e., independent given the values of the criterion on the objects), the optimality of *take the best*

(that searches cues in the order c_1, c_2, \dots, c_n) and *tallying* can be characterized as follows [KM00].

Result 3. For conditionally independent cues, *take the best* is optimal if and only if $o_i > \prod_{k>i}(o_k)$, where $o_i = v_i/(1 - v_i)$.

Result 4. For conditionally independent cues, *tallying* is optimal if and only if $v_i = v$.

Other Tasks

The paired comparison task is related to other tasks like deciding whether an object is larger than a certain threshold (classification) or judging how large the object is (estimation). Fast and frugal heuristics have been studied for these tasks as well and have been again found to perform well compared to standard benchmarks. For example, Katsikopoulos, Woike, and Brighton (2006) found that simple, fast and frugal trees can make more robust classifications than discriminant analysis and trees used in artificial intelligence [BFOS84]. Research on other decision tasks is reviewed in Gigerenzer (2004).

Fast and Frugal Heuristics for Engineering?

According to Moses (2004), one of the main goals of the study of engineering systems is to deal with changes that occur during the life cycle of these systems. Change can be dealt with actively by building flexibility into the system, that is, allowing the system to perform a number of functions. Change can also be dealt with passively by building a robust system, that is, a system that does not lose much of its performance when conditions vary. Both flexibility and robustness are necessary when there is uncertainty in the system. One might even say that uncertainty represents a chance for improvement in that it motivates flexibility and robustness (perhaps this is what Moses means when he talks about “viewing uncertainty as an opportunity”, 2004, p. 6). The final challenge is to combine these properties with transparency and usability so that the system will be accepted by its users.

Our research program was not designed to study engineering systems. Our results about the normative success of heuristics were obtained in decision tasks like comparing dropout rates in highschools. We did not study how engineering students and practitioners compare, say, two product designs. Of course, at a certain level of abstraction, these are very similar tasks, but we do not want to downplay the potential influence of context. Thus, we see our results as making a methodological suggestion about a new program of research. Below we argue that a fast-and-frugal-heuristics approach to making decisions in engineering systems may be helpful.

Fast and frugal heuristics tend to be robust. There are more results than those presented here, to this effect. For example, Brighton (2006) has pitted the heuristics against powerful machine learning methods (like Quinlan’s ID3 method) using the minimum description length as a criterion of robustness. He found that, in many cases, heuristics compressed the data more than the machine learning methods.

Heuristics implement, with their focus on external outcomes, a more practical intelligence than mathematical methods that target full internal rigor. Furthermore, heuristics are less ambitious than methods that try to work all the time; heuristics are problem-specific and information-structure-specific. A given heuristic may be applied successfully only to those comparisons, estimations, classifications, or choices with certain statistical properties (i.e., flat or very skewed distribution of cue validities). Taken together, however, heuristics cover a wide spectrum of decision tasks. The set of heuristics has been called the *adaptive toolbox* [GS01].

The adaptive toolbox is a flexible system for decision making: To build the heuristics in it, one combines different building blocks (rules for searching for information, e.g., by validity, with

rules for deciding based on the available information, e.g. use only one cue). The building blocks, themselves, are based on core psychological capacities (e.g., recognition). The toolbox allows the introduction of new heuristics by (i) combining existing building blocks in new ways or by (ii) creating new building blocks based on newly discovered capacities.

Finally, fast and frugal heuristics are transparent: They are easy to understand and apply (and, hence, are more acceptable). This is so because of two reasons. First, heuristics are expressed as clear and simple sequential algorithms (e.g., *take the best*). Second, the way they represent the information they use is consistent with people's cognitive representations (e.g., in *take the best* validities can be cast in terms of frequencies, not conditional probabilities; see also [HLHG00]). Perhaps for these reasons, some practitioners, like medical doctors, advocate the substitution of classical decision analysis with fast and frugal heuristics [EEER01, Nay01, KF00]. Just like some successful methods for engineering decision making — Pugh's (1981) concept selection — heuristics can be used to generate, rather than to suggest or impose, new possibilities. We want to encourage academics and practitioners to explore the potential of fast and frugal heuristics in engineering.

3.10 Chatting Penguins — The continuous structuring of both technical artefacts and community in open source software development.

Anika König, Jörg Strübing

Using the case of the Linux kernel development project as an example, this paper discusses the way how structure comes about and is being maintained among human actors. By analyzing the special mode of production of a central piece of open source software in a spatially distributed, weakly structured community of developers and users, we show that social and technical structures are both relevant sources and important results of a continual stream of interactions, embedded in which an ongoing process of structure-building takes place. Community-building and technology development are intertwined processes not organized by a central agency but constantly coordinated and balanced by processes of negotiation and interactional alignment among those actors who contribute to their joint endeavor. Here, the tension between local and global plus the differences between varying local perspectives result in a structure-building process for which universal optimality based on complete knowledge and overarching rationality cannot be achieved — neither for the social structure of the producing community nor for the technical structure of their product.

3.10.1 Introduction

Recently, research on innovation processes has gained new attention. Within the regime of knowledge societies first of all a lack of scientific knowledge has become obvious: How do innovations proceed? How do we gain new knowledge? How do we structure previous knowledge with respect to current problems? What is the impact of organization on knowledge production?

Within this line of research successful and supposedly new ways of both organizing knowledge and producing new solutions are of pivotal interest. The driving force here is the expectation that these new forms might contain structural features of knowledge organization and problem-solving not only suitable in their primary domain but transferrable to other fields in which the handling of knowledge and innovation is among the focal activities. With regard to structures, the central issue of this volume, investigating innovation practices allows to study structures “in the making.” While structures usually do not tell much about their genesis, and thus do not contain sufficient information about their politics and their legitimacy [Sta95], studying processes of structuring “in action” accounts for uncovering the motives, intentions, and perspectives that are the driving forces of structuring in social processes and which are responsible for the actual gestalt of the structures under scrutiny.

Mentioning motives, intention, politics within the broader context of research on structures and innovation directs us to a serious shift of perspective when it comes to the social sciences as compared to the approaches of engineering or mathematics: Instead of constructing problem solving algorithms to actually solve knowledge problems, in social sciences the focus is on the analysis of ongoing processes of structuring and problem-solving in innovation processes. This has important consequences:

First of all, under scrutiny are not our own ideas on how structuring and innovation processes should proceed. As social scientists we are primarily concerned neither with solving problems of the real world nor with modeling general solutions for solving problems of certain types. Instead, our subject matter is a social field that already is innovative and that already solves its problems — more or less efficient, successful or not, in more conventional or in more innovative ways. Our task here is to find out how they proceed, whether they are successful and, if so, why. Moreover,

people acting in that field interpret their and their co-actors' actions prior to and independent from our analysis. Thus, our analytic endeavors have to deal with both actions and those interpretations of actions that themselves again motivate new actions and interactions. Alfred Schütz [Sch53],³, coined the term "constructs of the second degree" to indicate the special quality of sociological interpretations of what happens in the social world.

For the investigation of structure-building innovation practices the field of new technologies is an especially promising place. But it is by no means self-evident that practices of technological innovation are at the same time performed in innovative manners. Technical or scientific innovations might as well be — and are maybe more often than not — developed in fairly conventional organizational settings. However, open source software development has gained special attention within the field of innovation research because the mode of innovation performed in this field is based on a number of core features that seem to indicate an innovative innovation process that takes place in an unconventional organizational setting.

Our aim in this paper is to investigate open source software development as a structure-building and innovation practice that is not only innovative but that at the same time — and at a certain level of abstraction — shows a number of structural features similar to those of certain mathematical and engineering strategies for solving problems of optimization under conditions of uncertainty, bounded rationality, and multi-criteriaity.

Basically, open source software are software products developed and distributed under certain types of licenses that deviate significantly from common copyright regulations of commercial software. Though there are a number of different copyright models in the field of open source software, the key feature is that the source code of this kind of software is public and might be used and modified by whoever desires to do so due to the so-called "open modification principle" [Lau04],⁶. While not for all open source software, so at least for those published under the "General Public License" (GPL) — like LINUX there is one precondition: Whoever modifies open software is obliged to publish the modified version under the same license.

While proprietary software is usually developed and maintained by formal organizations, the striking feature of open source software development is (at least most often) a minor emphasis on formalized organizational structures if not the complete absence of a formal organizational frame. Instead of a company hiring programmers to develop a software product, the field of open source software development is characterized by individuals or smaller groups starting a non-profit development project and inviting others to volunteer in participation.

The usage of expressions such as "most often" or "usually" signifies the large variety of different models within the open source scene. Indeed, we find a broad scale of different organizational settings not at least due to the fact that open source software projects tend to grow both in size and — consequently — in organizational and technical complexity. A project once initiated by a handful of close friends might expand to a worldwide network of participants. This paper focuses on the Linux Kernel project, one of the largest and oldest open source projects — and of course the most paradigmatic.

A serious amount of social science research has already been conducted on the phenomenon of open source programming. Research, however, that predominantly deals with economic and psychological aspects of open source. For economists, the impact of the very special copyright regulations in this domain has been a particular matter of interest [Kol99, Ray97]. Likewise, the question of why people would "invest" a larger amount of work into a product without being its "owner" and without receiving material rewards like a salary, has raised questions basically inspired by economic ways of thinking in terms of rational choice and individual profit [BR03, Lut04].

Less research effort, however, has been spent on the process of innovation management and community building within the various open source communities. The crucial issue here is that the

process of producing an innovative technological product takes place under the above mentioned very special conditions. Consequently, with the absence of a comprehensive formal organizational frame like a company that might organize the innovation process, there is also no formal and reliable hierarchical order within which tasks might be delegated and disobeying to orders might be sanctioned with definite consequences. And, most important, this absence of a formal hierarchy obscures the distribution of decision rights with respect to the acceptance or the rejection of problem definitions, suggestions, and patches delivered by participants.

Thus, the focus of our research endeavor is precisely the process of ongoing decision-making in the development process of the Linux kernel and — at the same time — the community-building effects interwoven with this continual process of negotiations concerning the best ways to improve the product under construction.

We claim that the Linux kernel community is not to be seen as a pre-established social structure within which the actual development work takes place. Instead, what is seen as a community is a current state of social relations resulting from those interactions programmers continuously undertake in order to jointly develop the Linux kernel software. The self-organization of the Linux kernel community is part and parcel of the working process that at first sight is only dedicated to bringing forth a technical artifact. At the same time, the development process itself has to be organized and structured in a continual process taking into account conflicting requirements, uncertainty about the other participants' proficiency, limited knowledge on alternative solutions, and so forth.

The aim of this paper is both to sketch the structural conditions of the community that further the innovation processes taking place in Linux kernel development and to discuss the implications this has for innovation and optimization as linked to structure building. Also, with reference to the framework of this book, we want to contribute to an extended understanding of structures that also incorporates the social sciences viewpoint into an interdisciplinary perspective of the subject.

The following section will give a short overview about the basic features of the Linux kernel project, while the third part of this paper discusses these features with respect to the question of self-organization and structure-building. This includes the processes of decision-making in technical matters that are closely intertwined with community-building. The final paragraph draws a number of conclusions with regard to the central issues of this volume, that is, multi-criteriality, bounded rationality and uncertainty.

3.10.2 The specifics of open source and the Linux-project

Software license policies

An important aspect of software is the license under which the computer program runs. Particularly nowadays where ownership rights and commodification are increasingly debated — not only in software developers' inner circles but also in society in general — software license policies are gaining increasing attention. However, this has not always been the case.

Historically, in the early days of computing (i.e. the 1950s), software was a free give-away, delivered by the hardware manufacturers in addition to the purchased hardware [Gra02],13,. It was a central feature of those — usually only rudimentary — computer programs that their source code was made available to everyone who wanted to read, modify, or enhance it. Hence, before software was demerged from hardware in the late 1960s and its source code often "closed", the users themselves usually adapted the software to their particular need [Gra02, Web04].

The separation of hard- and software — a process commonly known as unbundling — was a precondition for the emergence of software as a self-contained product and implemented in 1969 [Gra02], 203,276,. Due to this development the legal status of unbundled software had to be

clarified as well. Consequently and strongly influenced by official copyright and patent policies, a considerable number of different software licenses was generated by different parties over the following decades.

Software licenses can generally be divided into several groups. One way to classify software is the division between open source and proprietary software. As a commodified good, proprietary software comes without source code, that is, as bare binary machine code, not modifiable by its users: “Its use, redistribution or modification is prohibited, or requires you to ask for permission, or is restricted so much that you effectively can’t do it freely” [Fou01]. Open source software, on the other hand, comes together with its source code which makes it possible to modify it. Usually, open source software licenses allow free copying and distribution of the respective computer program and encourage people to contribute to its advancement.

In addition to the concepts of open source and proprietary software, there are multiple kinds of software licenses, such as for example freeware, shareware, public domain software, semi-free software, etc. [Fou01, Gra02],278,. Consequently, software licensing is subject to continuous debates and negotiations. The same holds for licensing within the context of open source software, since there are also a number of licenses that are all open source but differ with respect to some aspects. These are for example the BSD-licences (Berkeley Software Distribution) FreeBSD, NetBSD, and OpenBSD, and — most commonly used — the GNU General Public Licence (GNU GPL) which had been worked out by the GNU project (Gnu’s Not Unix) since 1984 [Fou01]. Although in the beginning Linux was released with an own license, since 1993 it also runs under the GPL.

The most crucial points in the GPL — the most prominent and probably also the most widespread license — are that the software may be copied and distributed in any medium, provided that each copy includes a copyright notice. In addition, the program may be modified, although this modification has to be documented thoroughly. The source code, though not necessary for the program’s runtime version, always has to be attached to the program, or, alternatively, it has to be indicated where the source code can be easily obtained. Finally, the license is spread to every computer program that is derived from GPL-software which is known as the “viral effect” of the GPL. Nevertheless, licensing under the GPL does not mean that the respective software has to be given away for free — on the contrary, the license explicitly states that “You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee” [Fou91].

Modes of organization in open source projects

Open source projects, even if they use the same licenses for their products, internally might be organized in very different ways. According to Kogut and Metiu, the governance structures, as they call it, are fundamental for the development of the software itself, since “organization by which work is delegated influences the product design” [KM01],page 257,. There are more “democratic” projects as well as hierarchical ones, and both kinds of projects in every single case are organized in a different way. The Apache project for example, one of the more hierarchically organized projects, is headed by a group of important developers. Linux, although there also is a group of developers in charge of the important decisions to be made, in the last resort is still subordinate to Linus Torvalds — the “benevolent dictator” (ibid.). Hence, this mode of organization can be considered as an increase in hierarchy as compared to e.g. the Apache project.

In order to better understand the organization modes of open source projects, based on an empirical investigation Crowston and Howison (2005) suggest to characterize them by their grade of centralization. They conclude, just like Kogut and Metiu, that the organization modes of the projects under scrutiny are extremely heterogeneous. The projects vary from highly centralized to

very decentralized, from “stars” with one central node to so-called “thickets” where every node is linked to every other one [CH05]. Hence, if linked to the above mentioned idea that some projects are hierarchically organized whilst others are more democratic, this would mean that the former are more centralized than the latter. With respect to the link between structure and innovation, this means that innovation in open source software development can take many shapes but one in order to be successful.

Regardless of the many differences in internal open source project organization, there are some structural similarities to be found. These similarities basically concern the roles developers can take in the organizational structure of the project. First of all, there are the maintainers. The role of a maintainer is usually awarded due to a developer’s accomplishment with regard to the advancement of the software. Maintainers have a high responsibility and many obligations on the one hand, but on the other this role also provides them with reputation and prestige and, if it comes down to it, with the last say concerning the part of source code they are responsible for. Typically, there are two ways to become a maintainer: first, maintainership is assigned on the basis of meritocracy. This means that who puts much time and effort into the project resulting in a quantitative and qualitative high output, in turn is awarded a responsible position. Second, whoever creates a fundamental and substantial piece of code, at the same time becomes the maintainer for this part or module of the software. Regardless of how they have obtained their position, it is the maintainers who finally decide which pieces of code are included into the program, they are central nodes in the communication and interaction processes, and they are the ones who are responsible for the adherence to the time schedule for the release of new versions of “their” module. However, as soon as the maintainer does not show the necessary commitment any more, the position is occupied by or assigned to someone else. Hence, the system is very flexible and fluid [Ett04],page 179ff.,

In addition, there are the contributing developers. Depending on the project and its size, sometimes this group of contributors is again divided into active, passive, and other kinds of developers. These distinctions are mainly made when the project has many members whose contributions differ profoundly in their quality and extent.

Finally, again depending on the particular project, there might be additional roles to be taken. For the Linux project there is, of course, Linus Torvalds — the originator of the first source code version of the Linux kernel. He is surrounded by his close confidants, the “trusted lieutenants”, who relieve him from having to deal with everything that is not absolutely central to the development of the Linux kernel (see fig 3.62). This mode of organization is by and large unique in the open source software world, very few other projects have a leadership position comparable to that of Linus Torvalds in the Linux project.

However, with respect to the organization modes in open source software development — here in Linux kernel programming — not only the programmers’ positions in the project organization are of interest, but also what additionally constitutes the structure of the community: a special form of communication that is independent from national borders, languages, and cultures — using the English language as the lingua franca. This kind of communication, usually in the semi-public kernel mailing list (everyone who desires to do so can enlist in this mailing list) or personally via ICQ, email, etc., is the basis for the development of the Linux kernel. Here, negotiations about great and small decisions with respect to all aspects of the program take place.

There is another feature of the project’s organization that must be pointed out: the particular importance of the single programmer. While the project would not work without the extremely dense communication (usually there are a few hundred emails posted to the kernel list per day), on another level the programmer is most likely to work alone, namely when programming or testing the software. The results — be it a piece of newly programmed code, the discovery of a bug in




Figure 3.62: Communication and cooperation between the members of the Linux kernel development community (L = Linus Torvalds; TL = Trusted Lieutenants; ● = Credited Developer; ○ = Developer/User; → ← = Direction of communication, patch exchange, etc.).

someone else's code, a bugfix, etc. — are returned to the community via the list or a maintainer.

Finally, what is of particular importance for the structure of the community and the structure of the program alike is version control and management. Here, Linux differs from most other open source projects since it uses proprietary software for its version control, a program called BitKeeper. Until 2002, Linus Torvalds had organized the project “manually”, i.e. without utilizing any computer program intended for this purpose. The excessive demands of this task brought him to finally implement BitKeeper for the management of the many different versions of the Linux software. However, this decision led to deep resentment in the Linux community since many of its members preferred an open source tool to be used for this task. Hence, a heated debate took place, primarily focusing on whether Torvald's decision was politically and morally correct. Nevertheless, it was finally accepted that BitKeeper was used instead of the popular open source version control system CVS (Concurrent Versions System). This incident nicely shows that in the end, even if great parts of the community do not completely agree with Torvald's decisions, he actually does have the final say. It is another question, however, to what extent Torvalds can practically exercise this final decision right. This will be dealt with in section 3.10.3.

Linux within the open source world

On the one hand, some general features can be observed in all open source projects: maintainers are responsible for different modules or versions of the program, communication is essential for the program development, and although the project cannot be sustained without the community, the work of the single developer is an essential aspect of the entire social and technical formation. Yet, on the other hand, within the open source world, Linux occupies a unique position. This can be ascribed to several different aspects:

First of all, the project is extraordinarily successful — in 2002, Grassmuck estimated the program to be installed on as many as 30 million computers [Gra02],page 229,. Today, many major companies and administration departments have “migrated” from proprietary software to Linux; amongst advantages such as the possibility to modify the code, the software is known for its good applicability in networks. In addition, commercial and non-commercial distributions of the program

(examples for commercial distributions are SuSE Linux and Red Hat, and for the non-commercial distributions Debian GNU/Linux) also make it possible for the “regular user” to install and use the program without extensive programming skills. Besides, even the commercial Linux distributions are still much cheaper than conventional proprietary software. There might be other open source programs that are as highly qualitative as Linux, but the availability of comprehensive distributions makes Linux particularly popular and, thus, successful and widespread.

Linked to the popularity of Linux is the size of the project. Not only is the program installed on tens of millions of computers all over the globe, but what is at least as remarkable is the size of the community contributing to its development. On the one hand, there are the kernel developers — the “cr??me de la cr??me” of Linux programmers — but they only constitute a part of the community since in addition there are innumerable programmers who in one way or another contribute to the project. This can be in form of a little patch, drivers for new hardware, etc. And finally, a huge part of the community are the passive contributors — the ones who only report bugs but do not write programs themselves.

Although the size of the project itself is already remarkable, what is even more outstanding — and this is what is of interest to us here — is that despite and not because of its size the project works efficiently.

3.10.3 Self-organization within the Linux kernel project

Self-organization and meritocracy: a contradiction?

The notion of self-organization in the context of social phenomena has to take into account a serious difference to self-organization with respect to natural or technical phenomena. Different from self-organization in e.g. molecule structures, clusters of microbes, or patterns of algorithms, self-organization of human actors has to deal with intentions, motivations, and interpretations. Human actors might not always — and most often do not at all — oversee the entire range of consequences their actions will bring about. However, when acting, they usually have “something in mind,” consciously or not they act towards certain goals and by doing so they employ their own interpretations of their fellow actors’ intentions, their “expected expectations” [Mea34].

Thus, a study of self-organization within the Linux community cannot treat actors as isolated entities jointly forming certain patterns which result from actor-independent mechanisms of, for example, mutation and selection. Self-organization in society basically means the existence of a certain social order under conditions of the absence of powerful actors equipped with universal knowledge concerting the social process in a desired direction.

Taking a look at the Linux kernel community we cannot refrain from recognizing the central role the originator of Linux, Linus Torvalds, is playing in the development of both the technical artefact and the social aggregation. Also, there obviously is a number of core actors, particularly the kernel maintainers, who do have a larger amount of power than other members of the community with respect to the most common and thus most important types of decisions to be taken within the project, for example rejecting or accepting a patch. These facts indicate some type of hierarchy present in the Linux kernel project. However, we would still claim that self-organization is a central feature of the Linux community’s social organization and of its innovation process management. Why is this so?

Despite the existence of powerful actors in the center of the development process, there also is a serious amount of power on the periphery: Different from a company or a bureaucracy, within the Linux community members are not members by contract but they are volunteers each and every time they contribute to the project. While it is true that Torvalds or his “trusted lieutenants” have the power to reject contributions as insufficient, ill-structured or faulty, they cannot, at the same

time, compel contributions from other actors. The Linux project depends on contributions from a larger community of skilful and motivated members since the project is by far too large to be developed by a narrow crew of programmers. Moreover, the quality of the program as a distinctive feature as compared to the Microsoft Windows operating system is based on a large network of competent user-developers. Also, acknowledging the membership of an actor is not a decision to be taken by central actors but it takes place in the course of communication within the public sphere of the project, for example in the mailing list: A formal boundary between community members and non-members does not exist. The more people are engaged in discussions within the list and contribute to the shared goal, the more they are considered to be members of the kernel programmers' community. E.C. Hughes has coined the term "going concern" for this type of dynamic social aggregations [Hug71]. Instead of being a matter of formal acts, inclusion or exclusion are a matter of developing a mutual understanding of the value an actors' contributions have for achieving the however roughly defined common goal. This implies that membership is not a discrete category but a gradual concept: Between a Linux user once sending a mail to the list and Torvalds as the center of the whole endeavour there is a continuum of membership intensities, organized by a continuous implicit and explicit evaluation of each member's contributions to the community's core activities.

Since much has already been written about Linux programmers' motivation to contribute to the project without being contracted and paid, we will not discuss this matter in detail. With respect to the question of self-organization, the important point is that the power to overrule decisions in every single case does not imply the same power for all cases to be decided. Contributions from users and independent developers are based on the actors' expectation to participate in a joint decision process. Since the Linux kernel community is based on meritocratic principles [Web04],page 180ff, participants tend to accept a larger influence of key actors — as long as their power and status is both justified by a fair amount of merits gained in the project and executed in a restrained manner. The notion of Torvalds being a "benevolent dictator" — as often cited in the Linux literature — signifies this attitude nicely: Since he is the originator of Linux, participants in the Linux kernel development would grant him a last word on every decision to take. However, a look at the Linux kernel mailing list shows that the frequency of Torvald's interference with cases to decide, at least within the list, is remarkably low.

As we have shown, besides personal preferences and habits there also is a functional explanation for this restrained use of directive rights by central actors within the Linux community: Not only are core actors like Torvalds unable to force contributions from other members, they are likewise in need to keep their contributors "in the mode" to participate in the programming endeavour. Although the functional perspective might hold analytically, it is, at the same time, not likely to be the actors' basic rationale within the Linux kernel community. Deeply rooted in the tradition of anti-commercial, anti-hierarchical movements the spirit of open source projects like Linux is still egalitarian and — in a grassroots manner — democratic.

With respect to self-organization, the key problem of open source projects lies in establishing a productive balance between central and peripheral clusters of power on the one hand (Torvalds, his "trusted lieutenants", and other maintainers) and the participation rights of "ordinary" members on the other. Essentially, this balance has to be produced and maintained continually throughout the project. In this respect, technical media (like BitKeeper or the Linux kernel mailing list), or informal agreed upon rules of conduct can be read as stabilizing features of the community, translating the needs of either group into the needs of the others, and, at the same time, representing a current state of the community's structure.

Distributedness and the need for effective communication

At this point we need to take a closer look at how the development work is actually done in the Linux kernel project. In the light of a world-wide distributed community of participants the organization of communication is pivotal for the progress of the project. In our case, geographically distributed participants at the same time means distributed work: Even tiniest pieces of software are manufactured from developers living on different continents, embedded in different cultures and speaking different languages. All the work performed under conditions so profoundly heterogeneous has to be brought together and integrated into one single software system.

Thus, the role of the Linux kernel mailing list cannot be overestimated. It is the media within which both the exchange of ideas and patches, and the negotiations over decisions concerning the program design basically take place. It is the virtual locale shared by and accessible to all actors already engaged in Linux development and also for those potential contributors still pondering about the idea of participating in the project. Besides Linux conferences or personal acquaintance with active developers, the Linux kernel mailing list is the central entrance to the Linux community. But more than this, it is the only community-wide arena where requests and claims can be made and negotiated.

A number of typical actions are likely to be taken by actors using the Linux kernel mailing list: Participants might

- complain about malfunctions of the kernel version
- suggest a new or modified functionality
- deliver a patch that supposedly adds or improves the kernel's functionality
- discuss advantages and disadvantages of suggestions or patches other presented by other participants
- discuss other matters like general design rationales, programming styles, or product politics

As the Linux kernel mailing list is the central communication forum for kernel related actions, virtually all topics relevant to the whole community are broadcasted through this channel. Of course, there is some more communication between single actors or smaller groups of actors not publicly observable, such as for example private mails or phone calls. However, all discussions and decisions relevant to both the product and the community usually pass the Linux kernel mailing list and are thus public.

Without a central steering committee, even the current agenda has to be negotiated through the list, though some of the general design perspectives will also be discussed on Linux kernel conferences. Different from formal organizations, negotiating the agenda in the Linux kernel mailing list is an ongoing process where consent is granted only until someone successfully opens up the discussion again. New topics can be brought up whenever someone feels the need to do so. Though obviously a general consent on some basic design rationales exists among most participants, various topics are repeatedly raised with a certain regularity, new suggestions are urged, and varying perspectives on current issues are presented.

The hitherto discussed aspects should have demonstrated how delicately the task of self-organization of the Linux kernel project as a community is embedded into the practical accomplishment of the collective development endeavour, while at the same time it is highly dependent on a communal communication forum. The important point is here that since the Linux kernel community lacks a central feature of traditional communities, that is, a spatial concentration allowing

for regular face-to-face interaction including diverse social activities, nearly all community-building has to be achieved through distant, work-related communication.

In an empirical analysis of thematic threads communicated through the Linux kernel mailing list, we found that while in the foreground negotiating over technical issues of certain patches or the overall design is processed, at the same time strategies are employed to socially grounding the interaction context. For example, a smoothing “social wrapping” [Hen99],page 62,like jokes, greetings or compliments often accompanies potentially conflict-laden contributions; likewise, actors jump in on harsher statements from other contributors, trying to moderate upcoming conflicts. True enough, similar moderations take place in other work-related types of communication too. However, in spatially non-distributed interactions social relations within the community can be maintained through other levels and forms of joint activity like having the proverbial chat at the coffee machine or going out for lunch. Thus, in the case of the Linux community social inclusion — as well as social exclusion — can be expressed near to exclusively by Linux-related communication through the mailing list. Repeated positive and supportive reaction to a contributor’s postings strengthens his or her position within the community — it produces reputation. However, this is tightly linked to actual contributions to the Linux kernel. The list is not meant to be used for small talk or for the exchange of compliments: “This list is not for the faint of heart” announces the introductory self-description of the list, demonstrating a strict dedication to the kernel as the only subject matter and to software development experience as the only acceptable habit (<http://www.ussg.iu.edu/hypermil/linux/kernel/info>; 27.07.05). That is to say, the “proof of the pudding” here is the submission and acceptance of patches improving the kernel software. Consequently, larger source code contributions and those that are valued as especially useful, innovative, and elegant increase its originator’s reputation. Within the mailing list, an ongoing communication about the acceptance of a patch gives both participants and passive observers (“lurkers”) a momentary impression of the contributor’s prestige. However, this is a fleeting impression, too weak to structure the community’s social order. Consequently, there are two more durable representations for the prestige-spending authorship of a piece of software that made its way into a kernel release: First, every patch is “signed-off” by its author, so the other developers working with the code would eventually come across the author’s name and electronic signature. Second, each Linux kernel release — like most other open source software — contains a “credit file” that lists the names of the contributing authors. Being mentioned there means that a developer’s achievements are made visible to a broader audience compared to the authoring signature of the patch. This is the reason why software companies when planning to hire new staff use to have a look at relevant credit files as a candidate’s recommendation.

Fluidity and durability in social order

A most suitable concept for understanding the structure-building processes within the Linux community is that of continual negotiations over status, reputation, roles, and legitimation. Negotiations not to be understood in terms of decision theory or market economy, that is, not necessarily as a formalized rational and strategic act, but as an often unintended and unnoticed process accompanying everyday interactions in every domain and resulting in a “negotiated order” [Str78]. This concept of social order holds that structures in societies are neither given nor fixed but subject to constantly negotiated changes. And, moreover, even maintaining an existing structure requires actors to confirm this order in their way of interacting. Without developers acting towards each other with a mutual understanding of a shared commitment to a common endeavour there would be no Linux kernel community. And without in their practical acting, ascribing certain developers a higher/lower status than others, there would be no “meritocracy”.

Of course interactive ascription and negotiation is only one, more fluid side of generating and maintaining social organizations as structures. The other, more durable side is the representation of a current state of affairs in the organizational memory of e.g. the layout of a technical infrastructure [Sta96]. In the case of Linux this is the organization of the mailing list that enables users to exchange their ideas, positions and information, and the structure of the source code repository (BitKeeper), supplying the participants with the latest version of the program and keeping them informed about changes. Another reification of structure is the modularization of the source code itself, representing a division of labour negotiated and agreed upon at some point back in the project's history. The definition of interfaces, the organization of certain functions in certain modules can be read as representations of the underlying social ordering of developers and users. Though information about structures is incorporated in these artefacts it should be stressed that these representations are not identical with the social structure itself. Since they are different from molecules, animals, or plants social actors have to interpret the artefacts as representations of structure. Structure comes into existence only through these mutual interpretations achieved in practical interactions. A simple example would be the structure of the source code: Each suggestion for implementing a new or altered functionality is both an interpretation of the current source codes structure and a potential challenge to it at the same time. The developers' interpretations, their collective negotiations about the suitability of interpretations, and their succeeding actions might — and often do — bring about a change of both the social and the technical structure. However, relying on artefacts as representations of structure implies that actors are not completely free in their actions: Whatever they interpret and negotiate, it is finally dependent on a proof in the "real world", it has to work both technically and socially.

Structuring of technical artefacts as a social process

While we have shed some light on processes of developing and maintaining a social organisation's structure, and its embeddedness in the process of technology development in the case of Linux, we have not yet had a closer look at the structuring of the technical artefact itself in the social process of joint distributed work in a heterogeneous community of developers and users. As shown in the second part of this paper, the basic process of open source software development takes place on two different levels: One is the semi-public negotiation process in the mailing list, the other is the process of writing and testing patches — a process largely performed by single developers.

The existence of these two different and to a great extent separated levels is of importance for the question of multi-criteriality and optimality. One strong motive for developers to participate in the Linux kernel project is the need to adopt the software to local conditions like e.g. a certain hardware configuration, special applications, or certain tasks to be performed. In our empirical analysis we encountered e.g. the case of a developer negotiating a redesign of two kernel modules in order to assemble a number of functionalities in one module while disposing of some other features. The local background for his demand was his work with Linux-based embedded systems like the ones used in digital watches or in MP3-players. The limited resources available in the application domains of embedded systems call for a rigid selection of only those functions needed for the special purposes addressed. The global view of maintaining and developing the general structure of the entire kernel had to take other issues into account, e.g. established design traditions and rationales that led to the modularization of the current release, or functional arguments valid for the full-fledged Linux in contrast to other more specialized usages.

It is right at hand that a universal overarching rationality suitable to manage the different perspectives maintained by the various local actors and groups is not available. Practical solutions have to be developed not only by taking into account different perspectives but also a broader number of additional criteria: While on the one hand the kernel as a whole needs to remain one

consistent piece of software (otherwise the joint development efforts would be split up in different projects, thereby suffering a loss of efficiency), it should, on the other, cover as many different purposes, applications, and technical constellations as possible. Furthermore, since the actual modularization represents not just a settled distribution of distinct program functionalities but at the same time a certain division of labor, including shared responsibilities and specialized knowledge domains, new suggestions have to be questioned as to whether they are in line with the established structure or if they can come up with reasons and arguments strong enough to pay off for the resulting organizational and social costs. Also, solutions will be judged with respect to the ratio of stability and innovativeness: System stability is a precondition for the kernel of an operating system, and by and large the old well-established, thoroughly tested system tends to be more stable and reliable than a system with newly implemented functions. However, an aging system cannot possibly keep track with the changing technological environment and with new application needs. Thus, a trade-off between risky innovativeness and the goal of system stability has to be made. A means taken by the Linux kernel community to handle this dilemmatic structure is to work with two separate kernel versions, one (with uneven version numbers) being experimental and thus less reliable while the other is a more settled and stable version (with even version numbers).

All these considerations form a — not always consciously addressed — background for the negotiations taking place on the kernel mailing list. Thus, whenever a patch is offered, choices have to be made: Take the patch as it is or suggest improvements? Insist on perfect and elegant problem-solutions or on a quick and dirty way to circumvent momentary ill-functions? Discuss only the patch on offer or take an initiative for a larger restructuring of the related part of the system? In processes of “interactional alignment” [Fuj87], the active participants try to fit a new proposal to the various constraints of the different technical, organizational and social levels. However, these processes are ubiquitous, they are not a privilege of open source development. Nevertheless, in cases like the Linux kernel project decisions have to be made in a semi-public participative process, that is, the other actors have to be convinced, they cannot simply be overruled.

In the case of the Linux kernel project the negotiation process is supported and stabilized by a rationalistic attitude typical for its participants (mostly software engineers and other “technology people”): In the light of a logic-based subject-matter like software they are convinced that this very subject matter itself is helpful in sorting the most suitable solutions out from ill-structured ones. “There is a strong sense among developers that disputes can and should be resolved “objectively” and that the best dispute resolution mechanism is simply to “let the code decide”” [Web04],page 164,. However, a look at the negotiations taking place in the Linux kernel mailing list reveals that this belief in the certainty of logic does not automatically lead to unequivocal decisions. While a larger number of cases can obviously be settled by checking the performance of a patch compared to the former state of the program, a still considerable amount of decisions lack such clear logical grounds. Though there is a limited number of clearly defined performance ratios that might be argued with, negotiations over patches rely on a broader set of — partly even contradictory — criteria, and some of them will be evaluated differently depending on the participants’ perspective.

One might ask: Why should Linux kernel developers care at all about other participants’ perspectives and requirements? Since all developers are free to pick a kernel version and to produce their own custom tailored versions by implementing locally desired features or structures, why should they bother about lengthy discussions? As already discussed, from the point of view of Linus Torvalds and his core team, integrating a heterogeneous and distributed crowd of developers is a necessary precondition to get the work done and the quality of the product thoroughly evaluated. But how about others, like e.g. the afore mentioned developer of embedded systems? He could have implemented the new module structure as required and have his system running. He could even get manufacturers of watches or — for that matter — intelligent refrigerators to use

his Linux version. His problem would be, however, that from this split off point on he would have to maintain his version without the established social and technical structure of the Linux kernel community. This invaluable advantage of a supportive knowledge base and work force of thousands of experienced developers is one of the pivotal incentives for staying within the community by negotiating towards compromises instead of producing individual or local versions.

Similar goals, different solutions

While most often negotiations within the list have to deal with acceptance, modification, or rejection of one certain patch offered by a single developer or a small team, in principle it is also possible that different patches dealing with the same problem are presented to the mailing list more or less simultaneously, offering different solutions for the same problem or adding the same functionality. Practically, however, this happens rarely. The unlikeliness of a parallel production of different solution to similar goals might be seen as an indicator for the strong inclusive effect of the Linux kernel mailing list and the BitKeeper source code repository. Contributors to the kernel source code are and need to be members of the mailing list plus they seem to monitor the discussion in the list thoroughly enough for not starting to develop a patch that others are already on the way to accomplish. Likewise, through the BitKeeper system all members have easy to use and real-time access to the most current source code version, thus new developments will not pass a developers attention unnoticed.

Despite mailing list and BitKeeper system, in principle it would be possible that two developers nevertheless start writing different code in order to solve the same problem independently from each other. This might happen because participants do not regularly or necessarily publish their intentions to develop a certain patch or feature. This is especially true for patches that solve problems claimed by the author of the patch. Here most often problem claim and suggested solution come in a package, ready-made for instant application. Though we did not perform a reliable check on this issue, a rough scan of the contents dealt with in the Linux kernel mailing list indicates to a large degree the absence of simultaneously developed competing patches.

Even forking, that is, the deliberate split off of one or more parties from a former joint development projects in order to develop different versions of the software (e.g. in cases of conflict over the design rationale or over the most elegant way to approach a problem) seems to be a rare phenomena in the Linux kernel community. Different from the afore discussed option of producing a custom tailored version of the kernel and thereby “leaving” the community, forking would not result in an irreversible loss of the community but for the kernel project it would produce the need to decide whether to integrate one or the other solution, none of them at all, or both. The possibility to accept two or more solutions for the same feature is most often limited by technical constraints. First of all, implementing alternative solutions for the same function results in a product more complex than (technically) required and thereby causing e.g. runtime problems and disproportionate complexity. Secondly, because of temporal or other resource limits especially in the kernel of an operating system a larger number of functions do not allow for dual or multiple structures at all. On the other hand there are aspects of the program that call for adaptability, and that is, for the users’ choice between different alternative functions already implemented as options. For Linux this is true for instance for the user interface.

Thus, unlike evolutionary adaptation processes in other domains, in Linux kernel development competition usually does not take place on the level of different developers’ solutions (competing patches) but instead on the level of competing arguments and evaluations judged by peers via the mailing list. Through the respective maintainer the community chooses not among various solutions but among the most convincing statements pro or contra the application of a certain solution offered by a contributor.

3.10.4 Conclusions

Though so far we have discussed social and technical innovation processes in the Linux kernel project, the aim of this paper is to only use this project as an example to illustrate a number of issues related to social structure-building. A number of key concepts with regard to structure have been mentioned in the introductory chapters of this volume: Multi-criteriality, optimality, bounded rationality, uncertainty. How then are these key concepts relevant and applicable if — as in our case here — the domains of structure-building are society and technology development instead of nature?

Analytically, we can recognize different types of structures in social phenomena on different levels of scale (micro, meso, macro; or — for that matter — interaction, organization, society), some of which are known to the actors and some not. Although social structures are obviously part of the adaptation of society and its parts to prevailing conditions, they are far from being optimal — as can be noticed in discussions over e.g. social inequality, urbanization, education systems, or globalization issues.

The case of the Linux kernel community shows aspects of structure on all these levels. It is part of a social movement and it represents a special form of globalized production based on a new type of organization combined with certain rules of interaction. In all these dimensions the structure of the Linux kernel community deals with structural conditions given in its ecology: Near-monopolistic market situations in the domain of operating systems; skill levels, traditions and fashions among programmers; the logical structure of programming languages; aspiration levels of different classes of users; availability and form of technical infrastructure, and so forth. Although the Linux kernel community could be understood as a response to these environmental conditions, its form and existence may not, however, be explained by these conditions. Quite obviously the social structure of this community is not determined by its ecology but there are elements of both coincidence and active choice besides those given conditions: Given its ecology, nobody could seriously have predicted the emergence of open source communities and their way of producing software. To generalize this issue: Human problem-solving does not result in completely determined and therefore predictable solutions — neither with regard to the subject-matter nor to the social organization of the issue of the problem-solving process employed. This inevitably results in forms of uncertainty to be taken into account in both the participants' actions and in the social scientists' analysis of these processes.

What is remarkable about social structures is the way how they are developed and maintained: Though structuring happens through active doing of humans, in most aspects this is achieved unconsciously. Actors perform various activities following goals or modes or routines, and in doing so they produce social structure as a by-product. As in our Linux case: In their deliberate efforts to produce an operating system in a way that is remarkably different from the way Microsoft produces, a growing number of developers and users by and by create and maintain a unique social structure which combines social community with aspects of effective work organization on a largely informal base.

Every effort to optimize the organizational structure of the community has to deal with a heterogeneous set of conditions resulting from serious differences in the various actors' perspectives: Contributing maintainers for instance might want to have immediate writing permission for the current source code version to thereby by-pass lengthy decisions processes in the mailing list. For other participants this would be a threatening prospect since they would lose their joint control over the kernel source code. Likewise, Torvalds and his core team of trusted lieutenants could easily be thought of as preferring direct and exclusive decision rights over design rationales and patch application in order to both speed things up and to make the kernel a unified whole.

Following either one of these perspectives would not lead to optimality, moreover, it is right at hand that the question of optimality cannot be judged from a universal rational point of view.

This is where multi-criteriality also comes into play. Whilst for some software developers the criterion of immediate individual writing permission is the foremost goal with respect to their contribution to the entire software system, for others joint control over the kernel source code is the most important criterion. In this case, in order to meliorate and optimize the version control system, the means of negotiation is employed as the best suitable means to reach the momentary stage of optimality. In this example this would be “as many rights as possible and as much community control as necessary”. It has to be kept in mind, however, that due to the characteristics of social systems, this stage of optimality at another instant might be something completely different.

The perspectives taken by the different types of participants in the Linux kernel community each represent a specific mix of requirements that would locally be judged as the optimal way of organizing the community and developing the source code, though the different local claims might conflict with each other. Thus, participants necessarily act based on a specifically bounded rationality. The final means that leads to the emergence of both the community as it is and the kernel in its current state is the continual stream of negotiations. No single actor or group of actors would have been able to establish the actual solution to the problem of organizing the kernel development. The knowledge structure emerging from these negotiations can be seen as an “intersection of independent lies” as the biologist Richard Levins puts it [Lev66],page 423,. The American sociologist Leigh Star explains the sociological essence of this notion: “???. Each local truth is partial and flawed; no a priori specification can encompass any global truth, but when (???) actors join local truths they create a robust emergent negotiated order” [Sta96],page 303,. The kind of cooperation discussed for the Linux case consists of such a joining of local truths — and this exactly is a pivotal dimension of inevitable uncertainty significant in our case: One can never know with certainty what the others know, want and will do.

In discussing matters of structure-building in society we have chosen the case of the Linux kernel project as an example. It might be asked, however, as to whether this example is representative or at least typical for all structure-building processes in societies? Of course it is not. There are various forms of how humans develop social structures, some are more overt and deliberate while others have the character of only happening occasionally and without the actors’ active knowing. The same holds for social structures themselves: Some are very formal, strict, and forcing, while others are more informal and weak. However, the focal point of our argument is that structure-building in the social sphere basically takes place in a continual negotiation process and that any social order constantly has to be maintained in interactions in order to exist — these points not only hold for the Linux kernel community but for society in general with respect to all its facets.

3.11 Optimization in Acoustics: Selected Examples

Peter Költzsch, Volker Bormann

The main objective of a technical system's optimization is, on first thought, an increase of its practical usability. Upon a second view the question arises, though, by what qualities this usability is defined. Which group of users has which demands? Are there additional, especially subjective, criteria that influence the decision process?

The field of acoustics is strongly prone to take subjective criteria into consideration as most, if not almost all acoustic systems strongly relate to the human auditive perception. And as soon as a perceptual factor plays an important rule in the assessment of a system, subjectivity will indispensably have to be dealt with in any optimizing procedure of respective systems. In some examples we will elucidate how this subjectivity may be take into account, partially by relating it to objectively measurable quantities, partially by just letting it govern a certain range of undecidedness or fuzziness.

3.11.1 "Good acoustics" of a concert hall or a opera house

Visitors of concerts or theaters expect, besides a good interpretation by the acting and playing artists, a technically good presentation of the performance at every location of the respective auditorium. They expect „good acoustics“ wherever they are seated. There is no precise definition of good acoustics in a room, though. This concept is rather described in fuzzy expressions, like audibility, acoustic properties of a room, acoustical intimacy, liveliness, reverberance, fullness of tone, sound volume, timbre, acoustic "warmth" of the room, spatial impression, auditory spaciousness, and acoustic envelopment of a listener ("a listener should be bathed in sound from all directions").

Room acoustic criteria have both subjective and objective aspects (AHNERT et al. 1993, ANDO 1985, 1988, ANDO et al. 1997, BERANEK 1992, 1994, 1996, 2003, FASOLD et al. 1984, 1987, 1998, KUTTRUFF 2000, REICHARDT 1979, SCHROEDER 1999). The subjective aspects relate to the perceived acoustic quality of a room, like the perception of acoustic spaciousness or clarity of the sound of the orchestra at music performances. The objective aspects are governed by the physical description of sound field parameters corresponding to the subjective quality assessment. Both kinds of criteria are also dependent on type and use of a room and on the type of an acoustic performance. Criteria for different kinds of music and for speech are to be distinguished.

The acoustic design of a room or a hall for the purpose of musical performances is usually part of an overall design. Thereby objective methods are used including objective criteria as well as subjective elements which accompany them. "Good acoustics" of a concert hall results from the interaction of subjectively assessed acoustic impressions in general and of subjective sound perceptions in detail as well as objectively sound field parameters, and that at points where listeners and musicians are placed.

A function sufficient for the goal "good acoustics" results from individual and from collective experiences with acoustically good or bad rooms. Furthermore such function of the goal can be derived from hearing tests in a room and from tests in a synthetic sound field in an anechoic room by mainly investigating separated sound effects.

Methods of computer simulation (with subsequent auralization) and investigations with acoustic models of a room, and applications of synthetic sound fields in an anechoic room can be used to calculating in advance or even to measuring the objective criteria in a not yet existing room.

As before mentioned at present no standard is at hand for the overall acoustic quality of a room or for "good acoustics". Especially a clear connection between objective factors which can be

measured to a subjective assessment is still partly unsolved. The acoustic literature offers a set of criteria for the assessment of the acoustic quality of rooms regarding musical performances (for example by FASOLD et al. 1987, 1993, 1998): - Spaciousness, spatial impression (temporal and spatial structure of a sound field), - Clarity (transparency), sound entry, absence of echo (temporal structure of a sound field), - Loudness, listening level, balance (dynamical structure of a sound field), - Timbre (spectral structure of a sound field).

The valence of these criteria approximately corresponds to the mentioned order. However it is absolutely dependent on genre and style, e.g. classical chamber music, romantic symphonic music, orchestral music, opera.

Importance	BARRON (1988)	SCHMIDT (1990)	BERANEK (1992)
1.	reverberation clarity	reverberation time	transparency time
2.	spaciousness	spatial loudness	impression
3.	loudness	loudness	spaciousness
4.	clarity timbre	clarity transparency	transparency
5.	intimacy	-	intimacy
6.	-	-	timbre (warmth)

Table 3.3: Comparison of evaluation of room acoustic quality characteristics of concert halls (FASOLD et al. 1993)

Descriptions of room acoustic quality characteristics (following FASOLD et al. 1998, AHNERT et al. 2002, MORFEY 2001, BERANEK 1996 etc.):

Reverberation time is the duration of audibility of reverberations. With it is connected the reverberation time which can be measured as time for a decay of sound energy of 60 dB after switching off the sound source. It is a rough quantity for the room acoustic quality of a concert hall, but it is not the only criterion.

Reverberance is the perception of reflected sound supplementary to direct sound which is not perceived as a repetition of direct sound.

Clarity, transparency is the possibility to perceive tones in a sequence of time (transparency in time) or in distribution of instruments (transparency in register).

Loudness, listening level is the decrease of sound pressure level in a room, in comparison with front places

Auditory spaciousness is the apparent enlargement of a source of sound (or orchestra) mainly in lateral direction to the audience in spite of its view (apparent source width)(BLAU 2001, 2002).

Auditory spatial impression is a being “wrapped in music”, an envelopment in music, that is the listener’s feeling, that the strength and directions of the reverant sound seems to arrive equally from all directions – forward, overhead, and behind (BERANEK 1996); a perception of a combination of all sources of sound with the environment of room and taking into account the audience therein.

Timbre is pitch feeling, sensation, warmth, brilliance, brightness.

Intimacy is the acoustic perception of the size of a hall, the feeling of proximity of a sound source, and interval in time between the arrival of direct sound and first reflections to the listener.

Definitions of the main acoustical criteria are given below, furthermore optimal ranges of large concert halls (FASOLD et al. 1998, BERANEK 1996, MORFEY 2001, AHNERT et al. 2002):

Reverberation time in seconds for decreasing the sound energy or the sound pressure level in a room from stop the source of its steady state level to one millionth or a level decay by 60 dB. The optimal range is 1.6 to 2.1 s. Early decay time is also applied as time in s for decreasing the sound pressure level in a room to the first 10 dB of the decay process. This time is most important for the perception of reverberance.

Reverberance measure is the ratio of reverberation energy in the time between 50 ms and ∞ (infinite) to that in the time between 0 and 50 ms. The optimal range is 0 to 4 dB for concert hall, 3 to 8 dB for symphonic music and -2 dB to $+4$ dB for musical theatres with optional use for concerts. A similar meaning has reverberation distance.

Clarity index, early-to-late sound index (fig. 6) is the ratio of early to late energy of sound energy arriving at a point in a room. For music is C80 valid. The limit between “early” and “late” is taken at 80 ms following the direct sound. The optimal range is -1 to $+3$ dB, depending strongly on the musical genre.

Early lateral energy fraction, lateral efficiency is the ratio of sound energy at a point in a room that arrives from the sides of a room to that which arrives from all directions, both within 80 ms. The optimal range is approximately 20

Spaciousness, spatial impression measure is defined as reinforcement of the spatial perception by reflections coming from the sides (LEHMANN 1974, SCHMIDT 1967, TRAUTMANN 1986, 1989, LEHNERT 1993). This index takes into account the spatial perception which is determined by reverberance and diffusivity (listeners envelopment) as also from the spaciousness which is caused by early lateral sound reflections (source broadening or apparent source width). For the acoustic quality of a concert hall are especially important the reflections from the side walls which are delayed to 25 to 80 ms compared to direct sound. The optimal range is $+1$ to $+7$ dB.

Interaural cross-correlation coefficient – IACC (MORFEY 2001) is a measure of the similarity between the signals received by the two ears and therefore a quantity of perception of a room. It is obtained by cross-correlating of the impulse response functions of a room which is measured at the positions of the two ears over a sample period. This is usually defined as the first 80 ms after the arrival of direct sound. With increasing influence of lateral reflections compared to direct signals of sound the IACC becomes less.

Loudness, sound level difference is the decay of a sound pressure level in a room. The level at any seat in a hall is compared with that in front of the podium (distance 5 to 10 m).

Sound entry (attack) is the perception of build-up time or the steepness of build-up of musical instruments. This is strongly influenced by technology of playing the instruments. It is characterized by the time of excitation. For shorter times the build-up of sound is steeper or the sound entry is harder.

Timbre (MORFEY 2001) is the tone colour or quality of musical sound. The term timbre can be applied to the sound of individual instruments or groups of instruments. It can also refer to the balance between low, middle and high frequencies in the musical sounds. Timbre is also an aggregate of attributes that allows a listener to distinguish a sound, in terms of subjective impression, from any other sound having the same loudness, pitch, and duration as well as the same direction of arrival.

Bass ratio (MORFEY 2001) is the ratio between the reverberation times at low frequencies (at octave centre frequencies 125 and 250 Hz) to reverberation times at middle frequencies (at octave centre frequencies at 500 and 1000 Hz) in a concert hall. The bass ratio is considered to relate to perceived warmth, that is the liveness or fullness of bass tones. Musicians sometimes describe a hall as „dark“ which has too prominent a bass or where high frequencies are largely attenuated (BERANEK 1996). The optimal range is 1.2 to 1.3.

Diffusivity is the local uniformity of a sound field distribution with respect to intensity and direction of incidence. It is also a temporal quantity for a statistical temporal distribution of a sound field. High temporal diffusivity means low appearance of harmonic natural frequencies.

Immediacy of response, attack (BERANEK 1996) is a feature from the musician's standpoint. A hall should give the performers the feeling that it responds immediately to a note. This term is related to the manner in which the first reflections from surfaces in the hall arrive back at the musician's ear. If the reflections occur too long after the note is sounded, they will be heard as an echo. If reflections are heard only from the nearby surrounding stage walls of the stage around him, the musician will fail to sense the acoustics of the hall altogether.

Results of assessments for two examples of well-known concert halls will be dealt with (FASOLD et al. 1987, 1993)

	Neues Gewandhaus Leipzig (1981)	Schauspielhaus Berlin (1985)
Hall:	fan-like plan	classical shoebox form
Volume:	21,000 m ³ , 1900 seats, 11 m ³ /seat	15,000 m ³ , 1430 — 1670 seats, 10.5 — 9 m ³ /seat

In these examples subjective assessments of acoustic quality were investigated with help of 50 to 60 test listeners at 5 selected representative seat areas. The assessments were required by the test listeners about the criteria: loudness, duration of reverberation, clarity, and spatial impression. These criteria seemed to be most important. For the assessments a scale of estimation was used. It was divided into 5 classes (middle class: without complaints, then: complaints are perceptible, and: complaints are disturbing; in both directions from low to high impressions). From the results conclusions were drawn of the acoustic quality of the room at individual seats as well as of the whole hall.

In fig. 1 the results of subjective tests are shown. The results are averaged over the seat areas.

Fig. 1: Selected results of subjective assessments of acoustic quality of two concert halls. The diagram shows the percentage distribution of these assessments for different acoustic qualities (FASOLD/STEPHENSON 1993)

As another example some results of optimization shall be dealt with for the new Semper Oper in Dresden. The optimization was carried out with subjective goal functions. The optimization was necessary because of the reconstruction of the opera-house after its destruction in World War II. It was reopened at February 13th, 1985. The large hall of this opera-house is of the type of a theatre

hall with semicircular plan and four balconies. Volume: 12,500 m³, 1,300 seats, 9.6 m³/seat. The hall can be used for performances of operas as well as for concerts. The investigations were carried out by Prof. Walter Reichardt and Prof. Wolfgang Kraak with their co-workers of the Institut für Technische Akustik der Technischen Universität Dresden in the time of 1970 and 1980 (KRAAK 1984).

On basis of existing documents a model of the historic hall of the Semper Oper was prepared on an scale 1:20, see fig. 2.

Fig. 2: View into the auditorium of the model (1:20) of the historic Semper Oper Dresden. Absorbing materials to imitate the audience are shown.(photo: R. DIETZEL, TU Dresden)

In this model the impulse responses of the room were measured as objective criteria. The measurements took place binaural with an artificial head. It was measured in the empty room and in the case of a opera-house with the fictitious audience. From the measurements the characteristics of sound quality were determined. These are reverberation quantity, clearness index C50, clarity index C80, spatial impression index, loudness, echo criteria, centre time, and reflection measure.

Hearing tests were carried out before opening the rebuilt Semper Oper. The tests included 86 persons in case of opera tests and 81 persons in case of concert tests. Overall 28 tests were carried out with different variants (spatial details, musical motives etc.). The test persons were 28

The criteria of assessments were dynamics (loudness, balance), temporal structure (reverberance, clarity (transparency), blend, echo), spatial structure (spatial perception, spaciousness), frequency composition (timbre, change in timbre), and acoustical overall perception. The quantity of hearing impression related to these characteristics had to be marked on a scale.

Following results were achieved: For concert cases the objective characteristics which were derived from impulse measurements are located on average in a optimal range. The differences between various seats in the audience and the differences between an occupied room and an empty room are relatively small. According to subjective assessments the majority of criteria on an average is found as optimal. In case of a concert it was found as acoustically extremely good.

For the opera case the objective tests showed the reverberation time as substantially dependent on the equipment of the stage. The other objective criteria are found partially in the optimal range and partially slightly off the optimal range. The subjective assessments showed in summary the following results: - Loudness, reverberation duration, mixing and spatial perception were estimated as optimal or as nearly optimal. - The transparency was assessed lower than for the concert case. - The speech intelligibility was assessed between good and satisfying. - The intelligibility of recitatives was assessed as better than good.

The comparison of acoustics of the old with the new Semper Oper was interesting (fig. 3). This comparison leads to the conclusions: The similarity in architecture of both opera houses had resulted in optimal values for spatial perception, mixing, and loudness. Therefore it can be stated the new opera house had attained the nearly former acoustic quality.

Fig. 3: View from the stage into the auditorium of the new Semper Oper (photo: R. Dietzel, TU Dresden)

The auditorium of the new opera house is slightly increased and therefore the compactness was slightly reduced. In consequence the reverberation time is increased. Now the slope of the parquet in the background and the view to the stage had improved the acoustic conditions. That's the reason for the good subjective assessments.

From now on the following investigation methods are used to designing acoustically a concert hall. This is necessary to follow the optimal acoustic criteria before a final design can be established.

Computer simulations as a method for acoustic design of a room The models of simulations are based on physical laws of geometrical room acoustics (ray acoustics). That means wavelength

are considered as short in comparison to characteristic dimensions of a room. There are in detail: mirror image source method, ray tracing models respective sound particle simulation methods (fig.4) including Monte-Carlo methods (VORLÄNDER/MECHEL 2002, STEPHENSON 2004)

Fig. 4: Computer simulation with a sound particle method (STEPHENSON in FASOLD et al. 1998)

The essential goal of computer simulations is to find out room impulse responses, fig. 5. From this the interesting acoustic criteria of a room will be calculated, e.g. clarity index C80 in fig. 6.

Fig. 5: Computer modelling in room acoustics At the left side is the acoustic system and on the right side the impulse response $h(t)$: from source to receivers (VORLÄNDER 2004)

Fig. 6: Results of a computer simulation: Distribution of the clarity index C80 in an concert hall with three balconies (VORLÄNDER 2004)

The computer simulation can be finished with an auralization. That means the audibility of a sound field can be realized at a seat in the auditorium of a virtual room by computational realization (MORFEY: "Auralization is the conversion of a digital waveform into audible form."). This auralization is carried out with help of calculated room impulse response, further with a "dry" (free of echo) recorded sound (e.g. music signals in anechoic rooms) as well as with outer ear transfer functions. Computer simulation with auralization thus enables us to hear into a yet not existing room!

Experimental model measuring technique for room acoustics The procedure is based on the similarity criterion of Helmholtz number $He = idem$ (KÖLTZSCH 2003). According to the reduction of geometry of a room (usual scale 1:20) the sound wave length are reduced in the same relationship. The frequency range is reciprocally transformed to the geometrical scale reduction (see example *Semper Oper* in fig. 2).

In a such model room the reflection and absorption properties of room surfaces must also be modelled. Sound sources are used e.g. shock excitation type spark transmitter with sound radiation characteristics which can be selected. The receiver is a model artificial head. It has two sound pressure receivers. Room impulse responses are determined in model measurements. This is the same as with computer simulations. Therefore the desired room acoustic characteristics are calculated (e.g. reverberation time, clarity index etc.). Furthermore the model offers the possibility of auralization (hearing into the model).

Model measurements are very expensive and time consuming. The advantage however consists in a clearness of the model. This is a crucial advantage in co-operation of acousticians and visually oriented architects. Diffraction and scattering processes in a room can be reproduced physically accurate which is a further advantage of model measuring techniques. These are the problem cases of a computer simulation. Details of configurations within a room can be varied fast and simply with model measurements. Examples are inclinations of walls or cover sections.

Synthetic sound field Natural sound fields can be reproduced in anechoic rooms (free-field rooms) with electro-acoustic equipments. For this purpose synthetic sound fields are generated with help of loudspeaker devices. An example shows fig. 7.

Fig. 7: Experimental set-up in an anechoic room to generating a synthetic sound field. The problem which is to be investigated in this case is the subjective investigation of the apparent source width of an orchestra (following M. BLAU in KÖLTZSCH (Ed.) 2003).

This research method optimizing room acoustics has the advantage to consist of the fact of systematically varying components of a single sound field. Furthermore the influence of variations of these components on room acoustics quality can be subjectively judged.

Optimal adjustment of hearing aids to rest-hearing ability of a deaf person Modern hearing aids have a large number of adjustable parameters. With this the loss of hearing can be individually compensated for with most different circumstances. Therefore the problem is to find out an optimal

parameter constellation “by hand” in order to maximally use the rest-hearing ability. Such a hearing aid adaptation is successively carried out in a “dialogue” between acoustician of hearing aids and hearing impaired persons. The optimization includes for developers and manufacturers of hearing aids absolute objective goal functions, among others: miniaturization of a hearing aid, minimization of energy consumption, expansion of the range of adjustable parameters, and adaptation to requirements of users.

Against a multi-criterial optimization by a hearing impaired person is primarily connected with subjective characteristics of his hearing impression. These are loudness, pleasantness (pleasant impression, being agreeable), timbre, distortions, speech intelligibility, noise, background and ambient noise. Furthermore non-acoustic, subjective criteria play frequently a crucial role, e.g. comfort of wearing, low maintenance, being user-friendly, service comfort, low purchase price, visibility in public.

For processes of adjustment is often used a so-called characteristics spider as a suitable tool. With it the multi-criteria optimization with subjective goal functions can be demonstrated, fig. 8.

Fig. 8: Characteristics spider for judgements of natural sound samples with an evaluation example (following G. FUDER in KÖLTZSCH (Ed.) 2003)

3.11.2 Optimization of the quality of speech communications

Optimization goal in speech communications is the highest possible degree of speech intelligibility. This has a considerable influence on successes into learning in schools, on aspects of occupational safety, and on the entire psychosocial conditions of an individual. Quantifiable and thus measurable objective criteria are the dimensions of a room, reverberation times, speaker-listeners distances, signal-to-noise ratio, type of background noise and its frequency spectrum, temporal structures, and information contents. Additionally measurable criteria are related to useful signals. These are speech levels, speech spectra, and speech velocities. Moreover for this optimization there are numerous subjective criteria. These cannot be measured but with them are connected important factors of influence: speech quality, speech melody (intonation, prosody), accents as well as dialects, possible hearing impairments, speakers and listeners knowledge of vocabulary of used languages in dependence on age, education, upbringing, interests, training and practice.

The optimization strongly is influenced by conditions as are types and places of speech presentations, are spoken texts as sentences, as single words, as numbers, as rhymes, and as sung texts. Places of speech presentations as churches, auditoriums, concert halls, or class or living rooms also are of influence.

Furthermore in context with speech presentations following visual subjective factors are of important influence for high speech intelligibility. These are visual contact to a speaker, gesture, facial expressions, mouth movements, and supporting graphical aids.

3.11.3 Acoustic optimization of modern floors in buildings (walking noise, impact sound)

Optimization of modern floors is highly multi-criterial. Primarily following goal functions are of use to manufacturers: profitability, competitiveness or marketability, acceptance of products by users, prices which can be achieved, quality control (guarantee, warranty), investment costs (production plants, buildings), material costs, production costs, and manpower costs.

Secondly for users the criteria of optimization with regard to acoustic characteristics are feeling at walking (elasticity and substructure of floors), walking noise (loudness, timbre, sound feeling),

and impact noise (excitation and propagation of structure-borne sound and its radiation into the surrounding air) (SARRADJ et al. 2000, 2001, 2004, 2005).

In a principle sketch are shown walking and impact noises (Fig. 9).

The process of optimization of these acoustic criteria including subjective judgements is very time consuming because of the dependence of walking noises from materials of a floor (elasticity, structure of sandwich-layers, mass), substructures, quality of installations, room acoustic conditions (especially reverberations), footwear, and walking habits (step rhythm, easy or sturdy steps).

Fig. 9: Walking noise and impact sound (principle sketch)(following E. SARRADJ 2004)

Subjective acoustic goal functions (quality criteria) of walking noise were fixed, concerning the potential buyers and people exposed to such unpleasant noise. This includes perceived loudness as objectively measured sound pressure level in average in a succession of pairs of steps, perceived timbre (pitch) as objectively measured amplitude spectrum of walking sound, and overall perception of walking noise with regard to pleasantness at sound exposure.

At hearing tests the walking sound of a person was measured for subjective assessments. 8 laminate floors from different manufacturers were investigated. They included different materials and layer structures. The before digitally recorded walking noises were offered in hearing tests to approximately 70 individuals. As method pair comparisons were chosen. Two probes of walking noises were offered over headphones. The test listeners had to mark on three scales which noise is quieter, is in timbre higher, and at least is more pleasant.

Results for all test persons were analysed with reference to sex and to age of a test group. These groups were divided into group 1 (under 30 years), group 2 (30 to 50 years), and group 3 (more than 50 years). An example of results shows fig. 10.

Fig. 10: Correlation of walking noise which is quieter and offers a more pleasant feeling (L-A) as well as a noise which is higher and offers a more pleasant feeling (H-A) (following R. EMMER, see also SARRADJ 2000, 2004).

The results lead to the conclusions: there is no walking noise on laminate floors which is ideal, the connection from perceived loudness and timbre to assessed pleasantness clearly is of subjective influence, and another influence is of sex and age group. Regarding the laminates judged as most unpleasant (loudest and duller walking noise) the assessments of all test persons were quite similar. On average of all test persons laminates were assessed as most pleasant which were quietest and made the highest timbre. Test persons of group 2 (30 to 50 years) gave the most stable assessments. With them the correlation between quiet and high to pleasant was the highest one. Male testers were more stable in their assessments. Female testers at least partly prefer louder and duller walking noises. Group 3 (more than 50 years) exhibits more weakly correlations (probable cause: their hearing impairments may have just started).

Conclusions: Loudness and pitch (timbre) as criteria which can be objectively measured are only a limited standard for pleasantness and subjective sound feeling. The loudness most influences pleasantness. Subjective assessments of timbre (deeper or higher pitch) vary individually. That means there is no distinct trend in the relationship of the measured criterion of timbre and the goal function of pleasantness.

Furthermore it is very interesting that the acoustic perceptions also depend on optical perceptions, in this case on décor (surface pattern) of laminate floor. Those laminates were assessed as acoustically most pleasant which also were assessed as more pleasant with reference to décor of laminate. This is also right regarding the order of pleasantness (SARRADJ et al. 2005).

Obviously an optimization of modern floors (laminate) cannot be carried out only with objective goal criteria. Of crucial importance also is the very complex subjective optimization criterion of pleasantness of walking noises.

3.11.4 Acoustic quality of loudspeakers

At least assessments of acoustic quality of loudspeakers are performed by subjective criteria. In hearing tests the overall subjective impression is assessed as subjective goal function of an acoustic optimization (KLIPPEL 1987, 1988, BECH 1992, OLIVE 2003).

Objective criteria for processes of optimization are sharpness, bass reinforcement, clarity, change in timbre, and distortions. Defect quantities are derived from these criteria as deviations from an optimum of each of them. The sum of all partial objective criteria is not identical to a subjectively perceived quality by the way of this optimization.

The acoustic qualification of test listeners is of great importance to this problem. Since this problem is generally relevant for optimizations with subjective goal functions here it should be dealt with the example of optimizations of loudspeakers. Trained experts as listeners are used in hearing tests of manufacturers. Such listeners represent subjective “measuring instruments” (KLIPPEL). The requirements to the listeners are hearing experience (experienced HiFi hearing, regular visit to concerts, musical activities by themselves). These experiences determine critically and discriminately the listeners conceptions of an ideal sound expression and the skill to being able to estimate a sound impression. Additionally test experience of trained listeners is an essential advantage. This includes familiarity with scaling methods and other conditions of hearing tests. In tests for loudspeakers as consumer products are preferred untrained persons (laymen). They can be seen as representatives for buyers. Demographic and socio-economic factors are of concern in these tests. Already the training process can influence the listener. Trained listeners are in general more critical than untrained listeners. For example they estimate acoustically bad loudspeakers worse than laymen. That means judgements of trained listeners can not be representative for unbiased and untrained listeners. This is valid for the large group of buyers and users. Expensive acoustic improvements of loudspeakers which were assessed by trained testers would often not be noticed by untrained listeners. Trained listeners are best in subjective tests with regard to reliability of assessments. At such tests high stress of testers has to be considered with regard to temporal and physical exhaustion. This may cause mental and physical fatigue and thus reduced attention and motivation.

Some information about a selected example follows now. In an extensive investigation four loudspeakers were tested (OLIVE 2003). For the purpose to estimate objective criteria for an optimization in an anechoic room sound signals were measured at conditions of a right angle to the diaphragm areas of the loudspeakers and in a hearing windows of horizontal ≈ 30 degrees and vertical ≈ 10 degrees. Additionally measured quantities were sound power, directivity index of sound radiation, and transfer functions.

In hearing tests an assessment scale was used for a perceived quality as a general impression (scale of assessment: 0 to 10 which means absolute rejection to total preference). The results distinctly correlate subjective judgements of the order of loudspeakers concerning their sound quality and the results of the measurements.

A significant result was the preference of some loudspeakers by assessments of untrained listeners. This is very similar to the assessments by trained listeners. Experts of these tests conclude hearing tests are the last crucial arbitrator of acoustic quality of loudspeakers.

Sound characteristics of loudspeakers should not only be assessed but also be described and scaled (KLIPPEL). Verbal terms are used for sound characteristics as volumes, sharpness, clarity, height or bass reinforcement, spatial feeling, change in timbre, bass clarity, and brightness. Objective assessments refer to descriptions of loudspeakers at listening conditions which were determined by measurements. The goal is to express the impression of listening by acoustic conditions which can be measured at a place of listening.

Subjective tests of loudspeakers can be represented with these results in a property space. Objective results of attributes of loudspeakers are combined in such a property space. The comparison of these two spaces shows whether the subjective test results can be explained with objective characteristics. As further result it can be shown whether subjective quality factors can be forecasted. Subjective loudspeaker tests should principally serve less to quality assessments according to opinions of manufacturers. These tests should serve more to subjective sound evaluations for users and buyers/customers.

A long path of research shall be gone to explain sound impressions of loudspeakers from physically measured quantities at conditions of listening places. This example seems to show the objectivation of subjective goal functions to be in a good way!

Acoustic quality of products and acoustic design (sound quality, sound design) The emitted sound is a quality feature of a set of products (JEKOSCH et al. 2004, 2005, BLAUERT et al. 1996, 1998, BODDEN et al. 2002, 2003). It is relevant to customers to purchasing a product. Manufacturer and buyer aim at and expect a sound which is a typical trade mark. Acoustic visiting card, "product sound" or corporate sounds are terms which are spoken of. Usually the following definition is used (BLAUERT/JEKOSCH 1996): Product-sound quality is a descriptor of the suitability of the sound attached to a product. It results from judgements upon the totality of auditory characteristics of the said sound – the judgements being performed with reference to the set of those desired features of the product which are apparent to the users in their actual cognitive, actional and emotional situation.

Nevertheless this problem is ambivalently judged. Some people prefer values in use in preference to structures of emitted sound. Others prefer quiet products. There are still others preferring louder noise and typical sounds of a product.

Further factors are important because the emitted sound has to have some content of information. This content is of concern to safety factors (e.g. in case of deviations from a normal machine noise or noise of an approaching vehicle as a warning signal to pedestrians). Noise can indicate operating conditions and necessities of maintenance or repair.

Emotional and very different meanings are linked to sounds in everyday life as is unquestionable. This depends on age, living standards, attitude towards life, hearing habits, and experiences. A consequence is the preference of some types or marks of products by different parts of a population. As an example some prefer extravagant and striking sound properties and others prefer particularly discreet sound qualities.

In actual cases it will be asked: What power of noise and which sound a mixer in a kitchen is allowed to have? What strength of noise, how loud, and which timbre a vacuum cleaner is allowed to have according to suction power or normal operating conditions? Which acoustic signal suggests overloading, necessity to empty the dust bag or even to repair? What sound of motor vehicle shall be established which is specific for a type of car (e.g. the so-called "Porsche sound")? What kind of sound shall be heard in a car coming from the windscreen wipers or the direction indicator? What kind of motor sound should be designed?

In an actual case these considerations lead to the question what really is the goal of optimization with so many subjective and objective criteria. Additionally in some cases must be taken into account limits of noise (e.g. noise with motor vehicles).

These problems of acoustic quality and acoustic design are rather similar to before mentioned ones. This leads to questions more in general: Can the chosen functions of subjective goals be definitely formulated? Is it possible to arrange a goal function in a form which can be measured? How large is the fuzziness of goal functions? How are to be selected testers for listening tests? How are to be trained testers?

Problems of acoustic qualities and acoustic design are characterized by the criteria (JEKOSCH et al. 2004, 2005) which limits of sound emissions are to be considered with respect to hearing damage and annoyance, which threshold values are to be considered as a safety factor to avoid risks or dangers and to judge operating conditions, which subjective preferences of noise are to be considered in very different situations of everyday life, and which quality of noise at a product and at production is to be considered according to requests of manufacturers and users.

To arrive at progresses of these problems empirical investigations are necessary in scientific research and in developments. Methods of evaluation must be achieved for acoustic product quality and acoustic design by means of field studies and laboratory tests. Goals of such investigations should be to derive rules for ingenious design of product sound. The aim is a code of practice for acoustical optimizations using subjective goal functions which are adapted to the products.

3.11.5 Subjective evaluations of acoustic qualities of musical instruments

For musical instruments the important acoustic criteria are pitch, response, timbre, projection etc. These can be derived from quantities of different measurements. Acoustic evaluations of musical instruments are traditionally conducted in a subjective manner up to now. But more and more a mix of subjective and objective criteria is used which complement each other. However in line with the trend objective criteria are applied which can be physically measured.

The statement of an expert is a generally accepted opinion: "An optimization is principally possible only if subjective criteria are described objectively. ... (But) the condition arrived at is still completely insufficient today." Therefore in this case the correlation between objective criteria and subjective evaluations is of crucial importance (BALTRUSCH 2003, GÄTJEN 1995, MEYER 1977, ZIEGENHALS 1995, 1996, 2000, 2002).

For subjective test procedures it is important to include judgements as to value by musicians and by listeners. However the test results of listeners and of musicians can deviate considerably from each other. As to the order of judgements musicians have a plain opinion: Goal of acoustic optimizations have to be fixed on ears of professional musicians. Objective criteria can only be used as a support to subjective impressions.

The discrepancies are well-known which arise from assessments of pitch and sound of guitars by listeners and musicians. Both use different criteria for their assessments. From the example with guitars of an acoustic optimization can be concluded to try more and more to correlate measured quantities and subjective judgements. A wide variety of instruments, large production output, and coercion to arrive at high acoustic quality gives the motivation to try this procedure.

Therefore in these multi-criteria optimization processes in most cases the poorly defined subjective criteria must definitely be described by physical quantities. Beforehand the subjective criteria should be assessed by trained listeners and professional musicians. Additionally, musical experts, manufacturers, and if necessary musical historians have to contribute. An influence of non-acoustic criteria takes part as well.

A further example of optimization is the quality judgement of bells (FLEISCHER 2000, HOUTSMA et al. 1987). According to expert opinions only subjective judgement is important. As an expert puts it: "The subjective judgement of a bell expert is above all results of physical measurements." (FLEISCHER 2003) In practice acoustic judgements on basis of subjective criteria are exclusively carried out mainly by special bell experts, church musicians, historians, but not by physicists or laymen. Both of the last would only be able to assess the sensory pleasantness of sound of a bell but poorly defined. The acoustic goal function of a bell is very subjective and of a multi-criteria nature. At present there is no objective procedure with a same result as subjective

evaluations. This is also true if in an objective procedure would be included a set of objective optimization criteria which can be measured.

Chapter 4

Transdisciplinary perspective

Transdisciplinary perspective

Bibliography

- [AEN79] A. Achilles, K.H. Elster, and R. Nehse. Bibliographie zur Vektoroptimierung. *Math. Op.forsch. Stat., Ser. Optim.* 10, (2), 1979.
- [AH05] A. Auger and N. Hansen. A restart CMA evolution strategy with increasing population size. In B. McKay et al., editors, *Proc. 2005 Congress on Evolutionary Computation (CEC'05), Edinburgh, Scotland*, volume 2, pages 1769–1776. IEEE Press, Piscataway NJ, 2005.
- [Aka73] H. Akaike. Information theory and an extension of the maximum likelihood principle. In *2nd International Symposium on information theory*, pages 267–281, Budapest, 1973. Akademiai Kiado, B.N. Petrov and F. Csaki.
- [AR93] N. Ambady and R. Rosenthal. Half a minute: predicting teacher evaluations from thin slices of nonverbal behavior and physical attractiveness. *Journal of Personality and Social Psychology*, 64:431–441, 1993.
- [Arn01] D.V. Arnold. Evolution strategies in noisy environments — A survey of existing work. In L. Kallel, B. Naudts, and A. Rogers, editors, *Theoretical Aspects of Evolutionary Computing*, Natural Computing, pages 239–249. Springer, Berlin, 2001.
- [ASK05] S. Ando, E. Suzuki, and S. Kobayashi. Sample-based crowding method for multimodal optimization in continuous domain. In B. McKay et al., editor, *Proc. 2005 Congress on Evolutionary Computation (CEC'05), Edinburgh, Scotland*, volume 2, pages 1867–1874. IEEE Press, Piscataway NJ, 2005.
- [Bäc96] Th. Bäck. *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, New York, 1996.
- [BB03] Th. Bartz-Beielstein. Experimental analysis of evolution strategies— Overview and comprehensive introduction. Interner Bericht des Sonderforschungsbereichs 531 Computational Intelligence CI–157/03, Universität Dortmund, Germany, 2003.
- [BB05] T. Bartz-Beielstein. *Experimental Research in Evolutionary Computation - The New Experimentalism*. Natural Computing Series. Springer, Berlin, 2005.
- [BB06] Th. Bartz-Beielstein. *Experimental Research in Evolutionary Computation—The New Experimentalism*. Springer, Berlin, 2006.
- [BBM04] Th. Bartz-Beielstein and S. Markon. Tuning search algorithms for real-world applications: A regression tree based approach. In G. W. Greenwood, editor, *Proc. 2004 Congress on Evolutionary Computation (CEC'04), Portland, OR*, volume 1, pages 1111–1118. IEEE Press, Piscataway NJ, 2004.

- [BCW04] Robert A. Burgelman, Clayton M. Christensen, and Stephen C. Wheelwright. *Strategic Management of Technology and Innovation*. 4th ed. McGraw-Hill, New York, 2004.
- [Bey95] H.-G. Beyer. Toward a theory of evolution strategies: On the benefit of sex – the $(\mu/\mu, \lambda)$ -theory. *Evolutionary Computation*, 3(1):81–111, 1995.
- [Bey00] H.-G. Beyer. Evolutionary algorithms in noisy environments: Theoretical issues and guidelines for practice. *CMAME (Computer methods in applied mechanics and engineering)*, 186:239–267, 2000.
- [BFM97] Th. Bäck, D. B. Fogel, and Z. Michalewicz, editors. *Handbook of Evolutionary Computation*. Oxford University Press, New York, and Institute of Physics Publ., Bristol, 1997.
- [BFOS84] L. Breiman, J. H. Friedman, R. A. Olshen, and C. Stone. *Classification and regression trees*. Chapman and Hall, New York, 1984.
- [BG97] Christof K. Biebricher and William C. Gardiner. Molecular evolution of RNA in vitro. *Biophys. Chem.*, 66:179–192, 1997.
- [BMK03] D. Büche, S. Müller, and P. Koumoutsakos. Self-adaptation for multi-objective evolutionray algorithms. In C. M. Fonseca, P. J. Fleming, E. Zitzler, K. Deb, and L. Thiele, editors, *Evolutionary Multi-Criterion Optimization, Second Int.'l Conf., (EMO 2003)*, number 2632 in LNCS, pages 267–281. Springer, Berlin, 2003.
- [Bol85] Béla Bollobás. *Random Graphs*. Academic Press, London, 1985.
- [BR85] R. Boyd and P. J. Richerson. *Culture and the evolutionary process*. University of Chicago Press, Chicago, 1985.
- [BR03] A. Bonaccorsi and C. Rossi. Why open source software can succeed. *Research Policy*, 32(7):1243–1258, 2003.
- [Bre62] H.J. Bremermann. Optimization through evolution and recombination. In M.C. Yovits, G.T. Jacobi, and G.D. Goldstein, editors, *Self-Organizing Systems*. Spartan Books, Washington DC, 1962.
- [Bru55] E. Brunswik. Representative design and probabilistic theory in a functional psychology. *Psychological Review*, 1955.
- [BS92] Thomas Bäck and Hans-Paul Schwefel. Evolutionary algorithms: Some very old strategies for optimization and adaptation. In D. Perret-Gallix, editor, *New Computing Techniques in Physics Research II*, pages 247–254. World Scientific, Singapore, 1992.
- [BS02a] Valerie Belton and Theodor J. Stewart. *Multiple Criteria Decision Analysis: An Integrated Approach*. Kluwer, Boston, 2002.
- [BS02b] H.-G. Beyer and H.-P. Schwefel. Evolution strategies – A comprehensive introduction. *Natural Computing*, 1:3–52, 2002.

- [BSS01] J. Branke, C. Schmidt, and H. Schmeck. Efficient fitness estimation in noisy environments. In L. Spector et al., editor, *Proc. of the Genetic and Evolutionary Computation Conference (GECCO'01)*, pages 243–250. Morgan Kaufmann, San Francisco, 2001.
- [BT93] J. R. Busemeyer and J. T. Townsend. Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 1993.
- [BT05] P.A.N. Bosman and D. Thierens. The naive MIDEA: A baseline multi-objective EA. In C.A. Coello Coello, A. Hernández Aguirre, and E. Zitzler, editors, *Proc. Evolutionary Multi-Criterion Optimization: Third Int'l Conference (EMO 2005)*, volume 3410 of *LNCS*, pages 428–442. Springer, Berlin, 2005.
- [BU88] Klaus Brockhoff and Christoph Urban. Die beeinflussung der entwicklungsdauer. *Zeitmanagement in Forschung und Entwicklung*, pages S. 1–42, 1988. Brockhoff, Klaus and Picot, Arnold and Urban, Christoph (Hrsg.).
- [CC05] J.L.A. Coello and C.A. Coello. MRMOGA: Parallel evolutionary multiobjective optimization using multiple resolutions. In D. Corne et al., editor, *Proc. 2005 IEEE Congress on Evolutionary Computation, (CEC 2005)*, volume 3, pages 2294–2301. IEEE Press, 2005.
- [CGG99] J. Czerlinski, G. Gigerenzer, and D. G. Goldstein. *How good are simple heuristics?* Oxford University Press, New York, 1999.
- [CH05] Kevin Crowston and James Howison. The social structure of free and open source software development. http://firstmonday.org/issues/issue10_2/crowston/index.html, 10.11 2005. 10/11/2005.
- [CM87] F. I. M. Craik and M. McDowd. Age differences in recall and recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 1987.
- [Coe02] C.A. Coello Coello. Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: A survey of the state of the art. *Computer Methods in Applied Mechanics and Engineering*, 191(11–12):1245–1287, 2002.
- [Coe06] C.A. Coello Coello. The EMOO repository: A resource for doing research in evolutionary multiobjective optimization. *IEEE Computational Intelligence Magazine*, 1(1):37–45, 2006.
- [Con96] J. Conlisk. Why bounded rationality? *Journal of Economic Literature*, 1996.
- [CRL03] Tim F. Cooper, Daniel E. Rozen, and Richard E. Lenski. Parallel chemages in gene expression after 20000 generation of evolution in *Escherichia coli*. *Proc. Natl. Acad. Sci. USA*, 100:1072–1077, 2003.
- [CS03] Y. Collette and P. Siarry. *Multiobjective Optimization. Principles and Case Studies*. Cost Engineering in Practice. Springer, Berlin, 2003.
- [CvVL02] C.A. Coello Coello, D.A. van Veldhuizen, and G.B. Lamont. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic Publishers, New York, 2002.

- [DA01] K. Dhami, M. and P. Ayton. Bailing and jailing the fast and frugal way. *Journal of Behavioral Decision Making*, 2001.
- [DAPM00a] K. Deb, S. Agarwal, A. Pratah, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. KanGAL report 200001, Indian Institute of Technology, Kanpur, India, 2000.
- [DAPM00b] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimisation: NSGA-II. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J.J. Merelo, and H.-P. Schwefel, editors, *Proc. of the 6th Int'l Conf. on Parallel Problem Solving from Nature - PPSN VI*, volume 1917 of *LNCIS*, pages 849–858. Springer, Berlin, 2000.
- [Dar59] Charles Darwin. *Of the Origin of Species by means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*, volume 811 of *Everyman's Library (Printing 1967, J.M.Dent & Sons, London)*. Murray, London, 1859.
- [Dar69] C. Darwin. *The autobiography of Charles Darwin*. Norton, New York, 1969. Original work published 1887.
- [De 75] K.A. De Jong. *An analysis of the behavior of a class of genetic adaptive systems*. PhD thesis, University of Michigan, 1975.
- [Dep03a] Department of Energy of the USA. Discussion of thermal comfort models. In *EnergyPlus Engineering Manual*, pages 171–190. 2003.
- [Dep03b] Department of Energy of the USA. EnergyPlus (e^+) Building Energy Simulation Software Home Page. <http://www.eere.energy.gov/buildings/energyplus>, (Stand Oktober 2003).
EnergyPlus is a elaborated simulation system for modeling the energetic flows and requirements in buildings, taking into account environmental conditions (solar irradiation, geographical orientation etc.) and usage dependent energetic inputs and losses. It is, in it's basic version, distributed free of charge.
- [dFIfTuW03] Bericht des Fraunhofer-Instituts für Techno-und Wirtschaftsmathematik, editor. *Radiotherapy — A Large Scale Multi-Criteria Programming.*, volume Nr. 43, Kaiserslautern, 2003. Intensity—Modulated.
- [DH01] M. K. Dhami and C. Harries. Fast and frugal versus regression models in human judgement. *Thinking and Reasoning*, 2001.
- [Dha03] M. K. Dhami. Psychological models of professional decision-making. *Psychological Science*, 2003.
- [DJW98] S. Droste, Th. Jansen, and I. Wegener. Perhaps Not a Free Lunch But At Least a Free Appetizer. Technical Report CI-45/98, Universit"at Dortmund, Dortmund, 1998.
- [DKS97] Andreas Drexl, Rainer Kolisch, and Arno Sprecher. Neuere entwicklungen in der projektplanung. *Zeitschrift für Betriebswirtschaftliche Forschung*, H. 2:95–102, 1997.
- [Dow01] S. Dowlatshahi. Product life cycle analysis: A goal programming approach. *Journal of the Operational Research Society*, Vol. 52(No. 11):1201–1214, 2001.

- [DP95] Diana De Pay. *Informationsmanagement von Innovationen*. Gabler, Wiesbaden, 1995.
- [ECL96] Santiago F. Elena, Vaughn S. Cooper, and Richard E. Lenski. Punctuated evolution caused by selection of rare beneficial mutations. *Science*, 272:1802–1804, 1996.
- [EEER01] G. Elwyn, A. Edwards, M. Eccles, and D. Rovner. Decision analysis in patient care, 2001.
- [Ehr00] M. Ehrgott. *Multicriteria Optimization*, volume 491 of *Lecture Notes in Economics and Mathematical Systems*. Springer, Berlin, 2000.
- [EL03] Santiago F. Elena and Richard E. Lenski. Evolution experiments with microorganisms: The dynamics and genetic bases of adaptation. *Nature Rev. Genetics*, 4:457–469, 2003.
- [ES90] Andrew D. Ellington and Jack W. Szostak. *In Vitro* selection of RNA molecules that bind specific ligands. *Nature*, 346:818–822, 1990.
- [ES02] A.E. Eiben and M. Schoenauer. Evolutionary computing. *Information Processing Letters*, 82(1):1–6, 2002.
- [ES03] A.E. Eiben and J.E. Smith. *Introduction to Evolutionary Computing*. Springer, Berlin, 2003.
- [Ett04] Matthias Ettrich. Koordination und kommunikation in open-source-projekten. In Robert A. Gehring and Bernd Lutterbeck, editors, *Open Source Jahrbuch 2004*, pages 179–192. Lehmanns Media, Berlin, 2004.
- [FF96] C.M. Fonseca and P.J. Fleming. On the performance assessment and comparison of stochastic multiobjective optimizers. In H.-M. Voigt, W.-Ebeling, I. Rechenberg, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature - PPSN IV*. Springer, Berlin, 1996.
- [FHM⁺06] Ricardo Flores, Carmen Hernández, A. Emilio Martínez de Alba, José-Antonio Daròs, and Francesco Di Serio. Viroids and viroid-host interaction. *Annu. Rev. Phytopathology*, 43:117–139, 2006.
- [Fix78] Wolfgang Fix. *Strukturuntersuchungen und Verfahren zur Zeitplanung bei GERT-Netzplänen*. Stuttgart: Hochschulverlag, Stuttgart, 1978.
- [Fog98] D.B. Fogel. *Evolutionary Computation: The Fossil Record*. Wiley–IEEE Press, New York, 1998.
- [Fou91] Free Software Foundation. General public licence. version 2. <http://www.fsf.org/licensing/licenses/gpl.html>, 1991. 11/22/2005.
- [Fou01] Free Software Foundation. Categories of free and non-free software. <http://www.gnu.org/philosophy/categories.html>, 11.22 2001. 11/22/2005.
- [FOW65] L.J. Fogel, A.J. Owens, and M.J. Walsh. Artificial intelligence through a simulation of evolution. In A. Callahan, M. Maxfield, and L.J. Fogel, editors, *Biophysics and Cybernetic Systems*. Spartan Books, Washington DC, 1965.

- [FS94] M. Forster and E. Sober. How to tell when simpler, more unified, and less ad hoc theories will provide more accurate predictions. *British Journal of the Philosophy of Science*, 1994.
- [FS98a] Walter Fontana and Peter Schuster. Continuity in evolution. On the nature of transitions. *Science*, 280:1451–1455, 1998. Also published as: Preprint No. 98-04-030, Santa Fe Institute, Santa Fe, NM 1998.
- [FS98b] Walter Fontana and Peter Schuster. Shaping space. The possible and the attainable in RNA genotype-phenotype mapping. *J.Theor.Biol.*, 194:491–515, 1998. Also published as: Preprint No. 97-11-081, Santa Fe Institute, Santa Fe, NM 1997.
- [FSS89] Walter Fontana, Wolfgang Schnabl, and Peter Schuster. Physical aspects of evolutionary optimization and adaptation. *Phys. Rev. A*, 40:3301–3321, 1989.
- [Fuj87] J. H. Fujimura. Constructing “do-able” problems in cancer research: Articulating alignment. *Social Studies of Science*, 17:257–293, 1987.
- [Ger04] Alexander Gerybadze. *Technologie- und Innovationsmanagement*. Vahlen, München, 2004.
- [GEZS06] Alexander E. Grobalenya, Luis Enjuanes, John Ziebuhr, and Eric J. Snijder. *Nidovirales*: Evolving the largest RNA virus genome. *Virus Res.*, 117:17–37, 2006.
- [GG02] D. G. Goldstein and G. Gigerenzer. Models of ecological rationality: The recognition heuristic. *Psychological Review*, 2002.
- [GGH⁺01] D. G. Goldstein, G. Gigerenzer, R. M. Hogarth, A. Kacelnik, Y. Kareev, and et al. Klein, G. Group report: Why and when do simple heuristics work? *Bounded rationality: The adaptive toolbox*, 2001. Cambridge, MA: MIT Press.
- [GGS⁺96a] Walter Grüner, Robert Giegerich, Dirk Strothmann, Christian Reidys, Jacqueline Weber, Ivo L. Hofacker, and Peter Schuster. Analysis of RNA sequence structure maps by exhaustive enumeration. II. Structures of neutral networks and shape space covering. *Mh.Chemie*, 127:375–389, 1996.
- [GGS⁺96b] Walter Grüner, Robert Giegerich, Dirk Strothmann, Christian Reidys, Jacqueline Weber, Ivo L. Hofacker, and Peter Schuster. Analysis of RNA sequence structure maps by exhaustive enumeration. I. Neutral networks. *Mh.Chemie*, 127:355–374, 1996.
- [GH97] D. Gigone and R. Hastie. The impact of information on small group choice. *Journal of Personality and Social Psychology*, 1997.
- [Gil76] Daniel T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comp. Phys.*, 22:403–434, 1976.
- [Gil77a] Daniel T. Gillespie. Concerning the validity of the stochastic approach to chemical kinetics. *J. Stat. Phys.*, 16:311–318, 1977.
- [Gil77b] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81:2340–2361, 1977.

- [GML⁺06] Kurt Grünberger, Ulrike Mückstein, Ulrike Langhammer, Andreas Svrcek-Seiler, Andreas Wernitznig, and Peter Schuster. RNA evolution *in silico*. Technical report, Institut für Theoretische Chemie, Universität Wien, 2006.
- [Gol89] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley, Reading, MA, USA, 1989.
- [GR87] D.E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodal function optimization. In *Proc. of the Second Int'l Conf. on Genetic Algorithms on Genetic Algorithms and Their Application*, pages 41–49. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 1987.
- [GR99] Monika Grötzner and Peter Roosen. Thermodynamic evaluation of chemical processes as a base of structure-optimizing process synthesis. In *Europ. Symp. on Computer-Aided Process Engineering*, Hungary, 1999.
- [Gra02] V. Grassmuck. *Freie Software. Zwischen Privat- und Gemeineigentum*. Bundeszentrale für politische Bildung, Bonn, 2002.
- [Gro05] Groundspeak Inc. Geocaching Portal. <http://www.geocaching.com>, September 2005. Introduction web site to the recreational use of the GPS technique for finding small hidden treasure chests. Being a relatively new occupation this idea seems to develop towards a kind of mass sports.
- [GS92] G. A. Geist and V. S. Sunderam. Network based concurrent computing on the PVM system. *Journal of Concurrency: Practice and Experience*, 4 (4):293 – 311, June 1992.
- [GS93] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 1993.
- [GS01] G. Gigerenzer and R. Selten. Bounded rationality: The adaptive toolbox. *Cambridge, MA: MIT Press*, 2001.
- [GS06] C. Grimme and K. Schmitt. Inside a predator-prey model for multi-objective optimization: A second study. In H.-G. Beyer et al., editor, *Proc. Genetic and Evolutionary Computation Conf. (GECCO 2006)*, Seattle WA, pages 707–714. ACM Press, New York, 2006.
- [GTtARG99] G. Gigerenzer, P. M. Todd, and the ABC Research Group. Simple heuristics that make us smart. *Oxford University Press*, 1999.
- [Gul50] H. Gulliksen. *Theory of mental tests*. Wiley, New York, 1950.
- [Gö95] Stefan Götze. *Die multikriterielle Entscheidungsfindung als Modell für die simultane Produktentwicklung*. Aachen: Shaker, Aachen, 1995.
- [Han99] T. Hanne. On the convergence of multiobjective evolutionary algorithms. *European Journal Of Operational Research*, 117(3):553–564, 1999.
- [Har04] F. L. Harrison. Advanced project management. In *A Structured Approach*, chapter 4th ed. Aldershot: Gower, 2004.

- [Has06] S. Hasse. *Taschenbuch der Geißerei-Praxis*. Schiele und Schön, Berlin, 2006.
- [Hen91] Ludwig Hennicke. *Wissensbasierte Erweiterung der Netzplantechnik*. Heidelberg, Berlin: Physica, Heidelberg, Berlin, 1991.
- [Hen99] J. Hendry. *Other people's worlds. An introduction to social anthropology*. Macmillan, Basingstoke, 1999.
- [Hen06] Al Hensel. Game of Life Homepage. <http://www.ibiblio.org/lifepatterns/>, (March 2006).
Original site for Conway's Game of Life. Contains a Java applet and some links to elaborate pattern providers.
- [Her97] M. Herdy. Evolutionary optimisation based on subjective selection — evolving blends of coffee. *Proceedings 5th European Congress on Intelligent Techniques and Soft Computing (EUFIT'97)*, pages 640–644, 1997.
- [HFS⁺94] Ivo L. Hofacker, Walter Fontana, Peter F. Stadler, L. Sebastian Bonhoeffer, Manfred Tacker, and Peter Schuster. Fast folding and comparison of RNA secondary structures. *Mh. Chemie*, 125:167–188, 1994.
- [HK05] R. M. Hogarth and N. Karelaia. Ignoring information in binary choice with continuous variables: When is less “more”? *Journal of Mathematical Psychology*, 2005.
- [HLHG00] U. Hoffrage, R. Lindsey, R. Hertwig, and G. Gigerenzer. *Medicine—communicating statistical information*, 2000. Science.
- [HM79] C.L. Hwang and A.S.M. Masud. *Multiple Objective Decision Making – Methods and Applications: A State-of-the-Art Survey*, volume 186 of *Lecture Notes in Economics and mathematical Systems*. Springer, Berlin, 1979.
- [Hof95] U. Hoffrage. *Zur Angemessenheit subjektiver Sicherheits-Urteile: Eine Exploration der Theorie der probabilistischen mentalen Modelle [The adequacy of subjective confidence judgments: Studies concerning the theory of probabilistic mental models]*. Unpublished doctoral dissertation. Austria, 1995.
- [Hof03] Ivo L. Hofacker. Vienna RNA secondary structure server. *Nucleic Acids Res.*, 31:3429–3431, 2003.
- [Hol73] J.H. Holland. Genetic algorithms and the optimal allocation of trials. *SIAM Journal of Computing*, 2(2):88–105, 1973.
- [HSF96] Martijn A. Huynen, Peter F. Stadler, and Walter Fontana. Smoothness within ruggedness: The role of neutrality in adaptation. *Proc. Natl. Acad. Sci. USA*, 93:397–401, 1996.
- [HT04] R. Hertwig and P. M. Todd. *More is not always better: The benefits of cognitive limits*. Reasoning and decision making: A handbook, 2004. Chichester: Wiley.
- [Hug71] E. C. Hughes. Going concerns: The study of american institutions. In E. C. Hughes, editor, *The sociological eye*, pages 52–64. Aldine, Chicago, 1971.

- [IL00] S. S. Iyengar and M. R. Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 2000.
- [Jäs01] Andres Jäschke. Artificial ribozymes and deoxyribozymes. *Curr. Op. Struct. Biol.*, 11:321–326, 2001.
- [Jel98] M. Jelasity. UEGO, an abstract niching technique for global optimization. In A. E. Eiben, Th. Bäck, M. Schoenauer, and H.-P. Schwefel, editors, *Proc. Parallel Problem Solving from Nature – PPSN V, Amsterdam*, pages 378–387. Springer, Berlin, 1998.
- [JGPP94] R. D. Jenison, S. C. Gill, A. Pardi, and B. Poliski. High-resolution molecular discrimination by RNA. *Science*, 263:1425–1429, 1994.
- [JR03] J. G. Johnson and M. Raab. Take the first: Option generation and resulting choices. *Organizational Behavior and Human Decision Processes*, 2003.
- [KC99] J. Knowles and D. Corne. The pareto archived evolution strategy: A new baseline algorithm for pareto multiobjective optimisation. In P.J. Angeline, Z. Michalewicz, M. Schoenauer, X. Yao, and A. Zalzal, editors, *Proc. Congress on Evolutionary Computation, (CEC'99)*, volume 1, pages 98–105. IEEE Press, Washington DC, 1999.
- [Ker05] Harold Kerzner. Project management: A systems approach to planning. In *Scheduling and Controlling Hoboken*. NJ: John Wiley & Sons., 2005.
- [KF00] K. V. Katsikopoulos and B. (in press) Fasolo. New tools for decision analysts. In *IEEE Transactions on Systems, Man and Cybernetics: Systems and Humans*, 0000.
- [KGC96] Elko J. Kleinschmidt, Horst Geschka, and Robert G. Cooper. *Erfolgsfaktor Markt: Kundenorientierte Produktinnovation*. Berlin et al.: Springer, Berlin et al., 1996.
- [Kim83] Motoo Kimura. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, 1983.
- [KLT03] T.G. Kolda, R.M. Lewis, and V.J. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Review*, 45(3):385–482, 2003.
- [Klu06] Sven Klussmann, editor. *The Aptamer Handbook. Functional Oligonucleotides and Their Applications*. Wiley-VCh Verlag, Weinheim, DE, 2006.
- [KM00] K. V. Katsikopoulos and L. (in press) Martignon. Naïve heuristics for paired comparisons: Some results on their relative accuracy. *Journal of Mathematical Psychology*, 0000.
- [KM01] B. Kogut and A. Metiu. Open-source software development and distributed innovation. *Oxford Review of Economic Policy*, 17:248–264, 2001.
- [KMM02] E. Kurz-Milcke and L. Martignon. *Modelling practices and “tradition”*. Plenum, New York, 2002. Model-based Reasoning: Science, Technology, Values.

- [KMS⁺05] KH Küfer, M Monz, A Scherrer, P Süß, , FV Alonso, AS Sultan, TR Bortfeld, and C Thieke. Multicriteria optimization in intensity modulated radiotherapy planning. *Berichte des ITWM*, 77, 2005. <http://www.itwm.fhg.de/zentral/download/berichte/bericht77.pdf>.
- [Kno02] J. Knowles. *Local-Search and Hybrid Evolutionary Algorithms for Pareto Optimization*. PhD thesis, The University of Reading, Reading, UK, 2002.
- [KOCZ93] G. A. Klein, J. Orasanu, R. Calderwood, and C. E. Zsombok. *Decision Making in Action: Models and Methods*, 1993. NJ: Ablex Publishing Corporation.
- [Kol99] P. Kollock. The economics of online cooperation: Gifts and public goods in cyberspace. In M. Smith and P. Kollock, editors, *Communities in Cyberspace*, pages 220–239. Routledge, London, 1999.
- [Koz92] J.R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press. Cambridge, MA, 1992.
- [KS05] N. Krasnogor and J.E. Smith. A tutorial for competent memetic algorithms: Model, taxonomy and design issues. *IEEE Transactions on Evolutionary Computation*, 5(9):474–488, 2005.
- [Lal01] K. N. Laland. Bounded rationality: The adaptive toolbox. In *Imitation, social learning, and preparedness as mechanisms of bounded rationality*. Cambridge, MA: MIT Press, 2001.
- [Lau04] A. M. St. Laurent. *Understanding Open Source & Free Software Licensing*. O’Reilly, Sebastopol, CA, 2004.
- [LBPC02] J.-P. Li, M.E. Balazs, G.T. Parks, and P.J. Clarkson. A species conserving genetic algorithm for multimodal function optimization. *Evolutionary Computation*, 10(3):207–234, 2002.
- [Lev66] R. Levins. The strategy of model building in population biology. *American Scientist*, 54:421–431, 1966.
- [LRS98] M. Laumanns, G. Rudolph, and H.-P. Schwefel. A spatial predator-prey approach to multi-objective optimization: A preliminary study. In A. E. Eiben, M. Schoenauer, and H.-P. Schwefel, editors, *Parallel Problem Solving From Nature — PPSN V*, pages 241–249, Amsterdam, Holland, 1998. Springer, Berlin.
- [Luc07] Klaus Lucas. *Molecular Models for Fluids*. Cambridge University Press, Cambridge, 2007.
- [Lut04] B. Luthiger. Alles aus spaß? zur motivation von open-source-entwicklern. In R. A. Gehring and B. Lutterbeck, editors, *Open Source Jahrbuch 2004*, pages 93–106. Lehmanns Media, Berlin, 2004.
- [Mah95] S.W. Mahfoud. *Niching Methods for Genetic Algorithms*. PhD thesis, University of Illinois at Urbana Champaign, 1995.
- [MD05] Martin G. Moehrle and Ewa Dönitz. *Die Notwendigkeit flexibler Mehrzielplanung von Innovationsprojekten*. Amelingmeyer, Jenny and Harland, Peter E. (Hrsg.), Wiesbaden: DUV, technologiemanagement & marketing edition, 2005.

- [Mea34] G. H. Mead. *Mind, self & society from the standpoint of a social behaviorist*. The University of Chicago Press, Chicago, Ill., 1934.
- [Meh05] J. Mehnen. *Mehrkriterielle Optimierverfahren für produktionstechnische prozesse*. Schriftenreihe des ISF. Hrsg.: K. Weinert. Vulkan Verlag, Essen, 2005.
- [MH02] L. Martignon and U. Hoffrage. lexicographic heuristics for paired comparison. theory and decision. *Fast, frugal and fit*, 2002.
- [Mie98] K. Miettinen. *Nonlinear Multiobjective Optimization*. Int. series in operations research and management science. Kluwer Academic Publishers, Boston, 1998.
- [Mie02] K. Miettinen. Interactive Nonlinear Multiobjective Procedures. In M. Ehrgott and X. Gandibleux , editors, *Multiple Criteria Optimization*, pages 227–276. Kluwer Academic Publishers, 2002.
- [MMBBH04] J. Mehnen, Th. Michelitsch, Th. Bartz-Beielstein, and N. Henkenjohann. Systematic Analyses of Multi-objective Evolutionary Algorithms Applied to Real-World Problems Using Statistical Design of Experiments. In R. Teti , editor, *Intelligent Computation in Manufacturing Engineering, 4th CIRP International Seminar on Intelligent Computation in Manufacturing Engineering, CIRP ICME '04*, pages 171–178, University of Naples, Italy, 2004.
- [MMLBB05] J. Mehnen, Th. Michelitsch, Chr. Lasarczyk, and Th. Bartz-Beielstein. Multi-Objective Evolutionary Design of Mold Temperature Control using DACE for Parameter Optimization. In H. Pfützner and E. Leiss, editors, *Short Paper Proceedings of the ISEM 2005, 12th Interdisciplinary Electromagnetic, Mechanic and Biomedical Problems*, volume L11-1 of *Vienna Magnetics Group Reports*, pages 464–465, Bad Gastein, 2005.
- [MMSK04] J. Mehnen, Th. Michelitsch, K. Schmitt, and T. Kohlen. pMOHypEA: Parallel evolutionary multiobjective optimization using hypergraphs. Technical Report of the Collaborative Research Centre 531 *Computational Intelligence CI-189/04*, University of Dortmund, 2004.
- [Mon01] D.C. Montgomery. *Design and Analysis of Experiments*. John Wiley, New York, 5 edition, 2001.
- [Mos89] P. Moscato. On Evolution, Search, Optimization, Genetic Algorithms and Martial Arts: Towards Memetic Algorithms. Technical Report Caltech Concurrent Computation Program, Report. 826, California Institute of Technology, Pasadena, CA, 1989.
- [MPS67] D. R. Mills, R. L. Peterson, and Sol Spiegelman. An extracellular Darwinian experiment with a self-duplicating nucleic acid molecule. *Proc. Natl. Acad. Sci. USA*, 58:217–224, 1967.
- [MR95] R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge University Press, New York, 1995.
- [MS96] Z. Michalewicz and M. Schoenauer. Evolutionary algorithms for constrained parameter optimization problems. *Evolutionary Computation*, 4(1):1–32, 1996.

- [MS05] Martin G. Moehrle and Wulf-Dieter Spilgies. Quality function deployment for product service systems. *Industrie Management*, 21. Jg(H. 3), 2005.
- [Nay01] C. D. Naylor. *Clinical decisions: from art to science and back again*. 358. The Lancet, 2001. Pachur, T., and Hertwig, R. (in press).
- [NM65] J.A. Nelder and R. Mead. A simplex method for function minimization. *Computer Journal*, 7(4):308–313, 1965.
- [OHMW02] T. Okuda, T. Hiroyasu, M. Miki, and S. Watanabe. DCMOGA: Distributed Cooperation model of Multi-Objective Genetic Algorithm. In *MPSN - II, The Second Workshop on Multiobjective Problem Solving from Nature*, Granada, 2002.
- [OW99] F. Oppacher and M. Wineberg. The shifting balance genetic algorithm: Improving the GA in a dynamic environment. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, editors, *Proc. Genetic and Evolutionary Computation Conf. (GECCO 1999)*, Orlando FL, volume 1, pages 504–510. Morgan Kaufmann, San Francisco, 1999.
- [Pap04] Paper presented at the Engineering Systems Symposium. *Foundational issues in engineering systems: A framing paper*, 29-31 March 2004. MIT, Cambridge, Mass.
- [Par96] V. Pareto. *Cours d’Economie Politique 1*. Lausanne, Rouge, 1896.
- [Par71] V. Pareto. *Manual of Political Economy*. Augustus M. Kelley (Original in French 1927), New York, 1971.
- [PNR06] M. Preuss, B. Naujoks, and G. Rudolph. Pareto set and EMOA behavior for simple multimodal multiobjective functions. In Th.Ph. Runarsson et al., editor, *Parallel Problem Solving from Nature (PPSN IX)*, volume 4193 of *LNCS*, pages 513–522. Springer, Berlin, 2006.
- [PSE05] M. Preuss, L. Schönemann, and M. Emmerich. Counteracting genetic drift and disruptive recombination in $(\mu \uparrow \lambda)$ -ea on multimodal fitness landscapes. In H.-G. Beyer, editor, *Proc. 2005 Conf. on Genetic and Evolutionary Computation, (GECCO 2005)*, pages 865–872. ACM Press, New York, 2005.
- [PSM⁺99] Dimitri Papadopoulos, Dominique Schneider, Jessica Meier-Eiss, Werner Arber, Richard E. Lenski, and Michel Blot. Genomic evolution during a 10,000-generation experiment with bacteria. *Proc. Natl. Acad. Sci.*, 96:3807–3812, 1999.
- [PVM] Home Page of PVM — Parallel Virtual Machine.
- [Ray97] Eric Raymond. Die kathedrale und der basar. das erfolgsgeheimnis von linux. *Linux-Magazin*, 8, 1997.
- [RdD97] Donna Cooper Richard de Dear, Gail Brager. Developing an adaptive model of thermal comfort and preference. Technical Report ASHRAE RP- 884, American Society of Heating, Refrigerating and Air Conditioning Engineers, Inc., http://aws.mq.edu.au/rp-884/RP884_Final_Report.pdf, March 1997. ASHRAE Report.

- [Rec71] I. Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. PhD thesis, Department of Process Engineering, Technical University of Berlin, Germany, 1971.
- [Rec73] I. Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog, Stuttgart, 1973.
- [Rec89] Ingo Rechenberg. Evolution strategy—nature’s way of optimization. In H. W. Bergmann, editor, *Optimization: Methods and Applications, Possibilities and Limitations*. Springer, Berlin, 1989.
- [Rec94] Ingo Rechenberg. *Evolutionsstrategie 94*. Stuttgart: Frommann-Holzboog, Stuttgart, 1994.
- [RK04] T. Reimer and K. Katsikopoulos. The use of recognition in group decision-making. *Cognitive Science*, 2004.
- [Roo99] P. Roosen. Concept and application of the evolutionary parameter optimization toolbox EPO. *ASME-Conference “Modeling and Simulation”*, 1999.
- [RSS97] Christian Reidys, Peter F. Stadler, and Peter Schuster. Generic properties of combinatorial maps. Neutral networks of RNA secondary structure. *Bull. Math. Biol.*, 59:339–397, 1997.
- [Rud98] G. Rudolph. On a multi-objective evolutionary algorithm and its convergence to the Pareto set. In D.B. Fogel, H.-P. Schwefel, Th. Bäck, and X. Yao, editors, *Proc. Fifth IEEE Conf. Evolutionary Computation (ICEC’98)*, Anchorage AK, pages 511–516. IEEE Press, Piscataway NJ, 1998.
- [Sac00] Christian Sachs. *Planung und Bewertung strategischer Investitionsprojekte auf Basis stochastischer Netzpläne*. Cottbus: Kovac, Cottbus, 2000.
- [SBA02] Günter Specht, Christoph Beckmann, and Jenny Amelingmeyer. *F&E-Management — Kompetenz im Innovationsmanagement*, volume 2. Aufl. Stuttgart: UTB, Stuttgart, 2002.
- [Sch53] A. Schütz. Common sense and scientific interpretation of human action. *Philosophy and phenomenological research*, 14(1):1–38, 1953.
- [Sch65] H.-P. Schwefel. Kybernetische Evolution als Strategie der experimentellen Forschung in der Strömungstechnik. Master’s thesis, Technical University of Berlin, Germany, 1965.
- [Sch75] H.-P. Schwefel. *Evolutionsstrategie und numerische Optimierung*. PhD thesis, Department of Process Engineering, Technical University of Berlin, Germany, 1975.
- [Sch77] H.-P. Schwefel. *Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie*, volume 26 of *Interdisciplinary Systems Research*. Birkhäuser, Basle, Switzerland, 1977.
- [Sch81] Hans-Paul Schwefel. *Numerical Optimization of Computer Models*. Wiley, Chichester, 1981.

- [Sch84] J.D. Schaffer. *Multiple Objective Optimization with Vector Evaluated Genetic Algorithms*. PhD thesis, Vanderbilt University, 1984.
- [Sch87] Hans-Paul Schwefel. Collective phenomena in evolutionary systems. In P. Checkland and I. Kiss, editors, *Problems of Constancy and Change – The Complementarity of Systems Approaches to Complexity, Proc. 31st Annual Meeting*, volume 2, pages 1025–1033. Int’l Soc. for General System Research, 1987.
- [Sch95a] H.-P. Schwefel. *Evolution and Optimum Seeking*. Sixth-Generation Computer Technology. Wiley Interscience, New York, 1995.
- [Sch95b] H.-P. Schwefel. *Evolution and Optimum Seeking*. Sixth-Generation Computer Technology Series. Wiley-Interscience, 1995.
- [Sch99] Heinz Joachim Schreyer. Untersuchungen zur zeitverkürzung des entwicklungsprozesses in der automobilindustrie mit der hilfe von gert-netzwerken. Technical report, Dissertation Technische Universität Cottbus, Cottbus, 1999.
- [Sch02] Rainer Schwarz. *Controlling-Systeme — Eine Einführung in Grundlagen, Komponenten und Methoden des Controlling*. Wiesbaden: Gabler, Wiesbaden, 2002.
- [Sch03] Peter Schuster. Molecular insight into the evolution of phenotypes. In James P. Crutchfield and Peter Schuster, editors, *Evolutionary Dynamics – Exploring the Interplay of Accident, Selection, Neutrality, and Function*, pages 163–215. Oxford University Press, New York, 2003.
- [Sch06] Peter Schuster. Prediction of RNA secondary structures: From theory to models and real molecules. *Reports on Progress in Physics*, 69:1419–1477, 2006.
- [Sci96] E. Sciubba. Numerical process- and plant simulation methods. *Chemical Engineering & Technology*, 19(2):170–184, 1996.
- [Sco02] A. Scott. Identifying and analysing dominant preferences in discrete choice experiments:. *Journal of Economic Psychology*, 2002.
- [SE97] G. Strunk and T. Ederhof. Machines for automated evolution experiments in vitro based on the serial tranfer concept. *Biophys. Chem.*, 66:193–202, 1997.
- [Sei00] R. Seindal. Gnu m4 development site. <http://www.seindal.dk/rene/gnu/>, (Stand Frühjahr 2000).
- [SF03] Peter F. Stadler and Christoph Flamm. Barrier trees on poset-valued landscapes. *Genet. Prog. Evol. Mach.*, 4:7–20, 2003.
- [SG02] Dieter Specht and Moehrl Martin G. *Lexikon Technologiemanagement*. Wiesbaden: Gabler, Wiesbaden, 2002.
- [Sha92] J. Shanteau. How much information does an expert use? is it relevant? *Acta Psychologica*, 1992.
- [Shi05] O.M. Shir. Niching in evolution strategies. In H.-G. Beyer, editor, *Proc. 2005 Conf. on Genetic and Evolutionary Computation, (GECCO 2005)*, pages 865–872, New York, 2005. ACM Press, New York.

- [Sim55] H. A. Simon. behavioral model of rational choice. *A. Quarterly Journal of Economics*, 1955.
- [Sim56] H. A. Simon. Rational choice and the structure of environments. *Psychological Review*, 1956.
- [SK00] Y. Sano and H. Kita. Optimization of Noisy Fitness Functions by Means of Genetic Algorithms using History of Search. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J.J. Merelo, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature (PPSN VI)*, volume 1917 of *LNCS*, pages 571–580. Springer, Berlin, 2000.
- [SKM⁺04] A Scherrer, KH Küfer, M Monz, FV Alonso, and TR Bortfeld. IMRT planning on adaptive volume structures - a significant advance in computational complexity. *Berichte des ITWM*, 60, 2004. <http://www.itwm.fhg.de/zentral/download/berichte/bericht60.pdf>.
- [SM01] W Schlegel and A Mahr. 3D Conformal Radiation Therapy - Multimedia Introduction to Methods and Techniques. Multimedia CD-ROM, Springer, 2001.
- [Spi71] Sol Spiegelman. An approach to the experimental analysis of precellular evolution. *Quart. Rev. Biophys.*, 4:213–253, 1971.
- [Spi06] Wulf-Dieter Spilgies. Realooptionen im produktinnovationsmanagement. Technical report, Dissertation Universität Bremen, Bremen, 2006.
- [SRB95] H.-P. Schwefel, G. Rudolph, and Th. Bäck. Contemporary evolution strategies. In F. Morán, A. Moreno, J.J. Merelo, and P. Chacón, editors, *Advances in Artificial Life – Proc. Third European Conf. Artificial Life (ECAL'95)*, pages 893–907. Springer, Berlin, 1995.
- [SSUZ03] F. Streichert, G. Stein, H. Ulmer, and A. Zell. A clustering based niching method for evolutionary algorithms. In E. Cantú-Paz, editor, *Proc. 2003 Conf. on Genetic and Evolutionary Computation, (GECCO 2003)*, pages 644–645. Springer, Berlin, 2003.
- [SSWF01] Bärbel R. M. Stadler, Peter F. Stadler, Günter P. Wagner, and Walter Fontana. The topology of the possible: Formal spaces underlying petterns of evolutionary change. *J. Theor. Biol.*, 213:241–274, 2001.
- [Sta95] S. L. Star. The politics of formal representation: Wizards, gurus, and organizational complexity. In S. L. Star, editor, *Ecologies of knowledge: Work and politics in science and technology*, pages 88–118. SUNY UP, Albany, NY, 1995.
- [Sta96] S. L. Star. Working together: Symbolic interactionism, activity theory and information systems. In Y. Engeström and D. Middleton, editors, *Cognition and communication at work*, pages 296–318. Cambridge University Press, Cambridge, 1996.
- [Sta98] P. Stagge. Averaging efficiently in the presence of noise. In A. Eiben, editor, *Parallel Problem Solving from Nature, PPSN V*, pages 188–197. Springer, Berlin, 1998.

- [Ste99] Steven Pinker. The Seven Wonders of the World. http://pinker.wjh.harvard.edu/articles/media/1999_06_07_mcgill.html, 7 June 1999. Steven Pinker's convocation address at McGill University.
- [Str78] A. L. Strauss. *Negotiations: Varieties, contexts, processes and social order*. Jossey-Bass, San Francisco, 1978.
- [SWJ98] M. Schonlau, W. Welch, and R. Jones. Global Versus Local Search in Constrained Optimization of Computer Models. *New development and applications in experimental design*, 34:11–25, 1998.
- [SWM⁺02] B. Schwartz, A. Ward, J. Monterosso, S. Lyubomirsky, K. White, and D. R. Lehman. Happiness is a matter of choice. *Journal of Personality and Social Psychology*, 2002.
- [SWN03] T.J. Santner, B.J. Williams, and W.I. Notz. *The Design and Analysis of Computer Experiments*. Springer Series in Statistics. Springer, New York, 2003.
- [TG90] Craig Tuerk and Larry Gold. Systematic evolution of ligands by exponential enrichment: RNA ligands to bacteriophage T4 DNA polymerase. *Science*, 249:505–510, 1990.
- [TGtArgip19] P. M. Todd, G. Gigerenzer, and the ABC research group (in press). Ecological rationality: Intelligence in the world. Technical report, Oxford University Press, New York, 19??
- [TLWK01] D. Thirumalai, Namkyung Lee, Sarah A. Woodson, and D. K. Klimov. Early events in RNA folding. *Annu. Rev. Phys. Chem.*, 52:751–762, 2001.
- [Tom05] M. Tomassini. *Spatially Structured Evolutionary Algorithms Artificial Evolution in Space and Time*. Natural Computing Series. Springer, Berlin, 2005.
- [Tve72] A. Tversky. Elimination by aspects: A theory of choice. *Psychological Review*, 1972.
- [Urs99] R.K. Ursem. Multinational evolutionary algorithms. In P.J. Angeline, editor, *Proc. of the Congress of Evolutionary Computation (CEC-99)*, volume 3, pages 1633–1640. IEEE Press, Piscataway, NJ, 1999.
- [vH05] Eric von Hippel. *Democratizing Innovation*. PhD thesis, Cambridge, MA: MIT Press, 2005.
- [Vie82] Dieter Viefhues. *Entscheidungskalküle zur Projektablauf— und —anpassungsplanung*. Frankfurt/Main: Lang, Frankfurt/Main, 1982.
- [VRMG04] Alexey G. Vitreschak, Dimitry A. Rodionov, Andrey A. Mironov, and Mikhail S. Gelfand. Riboswitches: the oldest mechanism for the regulation of gene expression? *Trends Gen.*, 20(1):44–50, 2004.
- [vV99] D.A. van Veldhuizen. *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*. PhD thesis, Air Force Institute of Technology, USA, 1999.
- [Vö71] Hubert Völzgen. *Stochastische Netzwerkverfahren und deren Anwendungen*. Berlin, New York: de Gruyter, Berlin, New York, 1971.

- [WA] Wikipedia-Article. Phylogenetics. <http://en.wikipedia.org/wiki/Phylogenetics>. Explanation of the concept of phylogenesis, i.e. the systematic order of ancestral relationships of biological species.
- [WB03] Wade C. Winkler and Ronald R. Breaker. Genetic control by metabolite-binding riboswitches. *Chembiochem*, 4(10):1024–1032, 2003.
- [Web04] S. Weber. *The success of open source*. Harvard University Press, Cambridge, MA, 2004.
- [WGW72] H. Wille, K. Gewalt, and H. D. Weber. *Netzplantechnik - Methoden zur Planung und Überwachung von Projekten*, volume Band 1: Zeitplanung of 3. Aufl. Muenchen: Oldenbourg, München, Oldenbourg, 1972.
- [Wie03] Wieder GmbH. Werkzeugkühlung erhält Impulse. *MaschinenMarkt, MM Plastics*, 5, 2003.
- [Wika] Wikipedia article. Civilization (computer game). http://en.wikipedia.org/wiki/Civilization_computer_game. Short introduction in a series of computer games called 'Civilization', with a discussion of their features and game systematics.
- [Wikb] Wikipedia article. Marsupial. <http://en.wikipedia.org/wiki/Marsupial>.
- [Wikc] Wikipedia article. Teleology. <http://en.wikipedia.org/wiki/Teleology>. Explanation of various philosophical concepts of an assumed targeted development, mostly of nature.
- [Wik06a] Wikipedia, <http://en.wikipedia.org/wiki/Kriging>. *Kriging*, July 2006. Wikipedia synopsis of the method and history of the Kriging interpolation technique.
- [Wik06b] Wikipedia article, <http://en.wikipedia.org/wiki/Prolog>. *Prolog, the programming language*, July 2006. Short overview about the concepts of the computer programming language Prolog.
- [Win72] A. Winfree. *Science*, 175:634–636, 1972. Experimental realisation of a Belousov-Zhabotinskii reaction.
- [WJ99] C. Weihs and J. Jessenberger. *Statistische Methoden zur Qualitätssicherung und -optimierung in der Industrie*. Wiley-VCH, Weinheim, 1999.
- [WM97a] D. H. Wolpert and W. G. Macready. No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.
- [WM97b] D.H. Wolpert and W.G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.
- [WMM⁺04] K. Weinert, J. Mehnen, Th. Michelitsch, K. Schmitt, and Th. Bartz-Beielstein. A Multiobjective Approach to Optimize Temperature Control Systems of Molding Tools. *Production Engineering ? Research and Development, Annals of the German Academic Society for Production Engineering (WGP)*, XI(1), 2004.

- [Wri95] M.H. Wright. Direct search methods: Once scorned, now respectable. In *Proc. 1995 Dundee Biennial Conf. in Numerical Analysis*, volume 344 of *Pitman Res. Notes Math. Ser.*, pages 191–208. CRC Press, Boca Raton, FL, 1995.
- [WS99] David S. Wilson and Jack W. Szostak. In vitro selection of functional nucleic acids. *Annu. Rev. Biochem.*, 68:611–647, 1999.
- [WWB05] H.A. Wichman, J. Wichman, and J.J. Bull. Adaptive molecular evolution for 13,000 phage generations: A possible arms race. *Genetics*, 170:19–31, 2005.
- [ZLB03] E. Zitzler, M. Laumanns, and S. Bleuler. A tutorial on evolutionary multiobjective optimization. In *Workshop on Multiple Objective Metaheuristics (MOMH 2002)*. Springer, Berlin, 2003.
- [ZLT01a] E. Zitzler, M. Laumanns, and L. Thiele. SPEA2: Improving the Strength Pareto Evolutionary Algorithm. Technical Report 103, Computer Engineering and Communication Networks Lab (TIK), Swiss Federal Institute of Technology (ETH), Gloriastrasse 35, CH-8092 Zürich, 2001.
- [ZLT01b] E. Zitzler, M. Laumanns, and L. Thiele. SPEA2: Improving the strength pareto evolutionary algorithm for multiobjective optimization. In K.C. Giannakoglou, D.T. Tsahalis, J. Periaux, K.D. Papailiou, and T. Fogarty, editors, *Evolutionary Methods for Design, Optimization and Control with Applications to Industrial Problems*, pages 1–6. International Center for Numerical Methods in Engineering(CIMNE), Barcelona, 2001.
- [ZS81] Michael Zuker and Patrick Stiegler. Optimal computer folding of large RNA sequences using thermodynamics and auxiliary information. *Nucleic Acids Res.*, 9:133–148, 1981.
- [Zuk89] Michael Zuker. On finding all suboptimal foldings of an RNA molecule. *Science*, 244:48–52, 1989.
- [Zö97] O. Zöllner. Optimierte Werkzeugtemperierung. Technical Report ATI 1104, Bayer, Plastics Business Group, 1997.

Appendix A

Contributing Scientists

Half a page per person. (Optional, should be discussed)