Nada är en gemensam institution mellan
Kungliga Tekniska högskolan och Stockholms universitet.

A Gaming Perspective on Command and Control

JOEL BRYNIELSSON

Avhandling som med tillstånd av Kungliga Tekniska högskolan
framlägges till offentlig granskning för avläggande av teknologie doktorsexamen
torsdagen den 15 juni 2006 kl 14.00
i sal E3, Osquars backe 14,
Kungliga Tekniska högskolan, Stockholm.

TRITA-CSC-A 2006:07
ISSN 1653-5723
ISRN KTH/CSC/A--06/07--SE
ISBN 91-7178-365-2

Abstract

In emergency management and in military operations, command and control comprises the collection of functions, systems and staff personnel that one or several executives draw on to arrive at decisions and seeing that these decisions are carried out. The large amount of available information coupled with modern computers and computer networks brings along the potential for making well-informed and quick decisions. Hence, decision-making is a central aspect in command and control, emphasizing an obvious need for development of adequate decision-supporting tools to be used in command and control centers. However, command and control takes place in a versatile environment, including both humans and artifacts, making the design of useful computer tools both challenging and multi-faceted.

This thesis deals with preparatory action in command and control settings with a focus on the strategic properties of a situation, i.e., to aid commanders in their operational planning activities with the utmost goal of ensuring that strategic interaction occurs under the most favorable circumstances possible. The thesis highlights and investigates the common features of interaction by approaching them broadly using a gaming perspective, taking into account various forms of strategic interaction in command and control. This governing idea, the command and control gaming perspective, is considered an overall contribution of the thesis.

Taking the gaming perspective, it turns out that the area ought to be approached from several research directions. In particular, the persistent gap between theory and applications can be bridged by approaching the command and control gaming perspective using both an applied and a theoretical research direction. On the one hand, the area of game theory in conjunction with research findings stemming from artificial intelligence need to be modified to be of use in applied command and control settings. On the other hand, existing games and simulations need to be adapted further to take theoretical game models into account.

Results include the following points: (1) classification of information with proposed measurements for a piece of information's precision, fitness for purpose and expected benefit, (2) identification of decision help and decision analysis as the two main directions for development of computerized tools in support of command and control, (3) development and implementation of a rule based algorithm for map-based decision analysis, (4) construction of an open source generic simulation environment to support command and control microworld research, (5) development of a generic tool for prediction of forthcoming troop movements using an algorithm stemming from particle filtering, (6) a non-linear multi-attribute utility function intended to take prevailing cognitive decision-making models into account, and (7) a framework based on game theory and influence diagrams to be used for command and control situation awareness enhancements. Field evaluations in cooperation with military commanders as well as game-theoretic computer experiments are presented in support of the results.

Keywords: command and control, decision-making, situation awareness, data fusion, simulation, gaming, experimentation, microworld research, graphical modeling, game theory, rationality

Contents

Chapter 1

Introduction

Natural and man-made events causing harm to people's life, property, living conditions or industry, are examples of events that exploit various kinds of society vulnerabilities. Large-enough events often result in disasters that affect human societies, ecosystems and environment negatively for long periods of time, requiring various degrees of resilience. Command and control (C2) involves decision-making and decision execution to reduce the need for resilience in these kinds of situations. A C2 system supports decision-making and decision execution in C2.

The topic for this thesis is C2 decision-making, i.e., the act of coming up with the, in some sense, "best" decision in a C2 situation. More precisely, for the most part we will be interested in decision-making characterized by varying amounts of conflict, i.e., situations influenced by several opposing actors. Man-made disasters typically involve a certain degree of conflict where opponents try to outperform each other, but natural disasters may also be seen as conflict situations against the laws of nature. The viewpoint taken in this thesis is therefore to treat natural disasters as similar to man-made disasters.

1.1 The Gaming Perspective

Situations handled by C2 systems develop, among other things, according to actions undertaken by opposing decision-makers. Also, C2 situations develop according to a number of other more or less uncertain factors. This combination of strategic interaction and situation complexity gives rise to the game arena targeted by the results presented in this thesis. That is, a C2 situation is neither a pure game with precise rules, nor is it a situation that can be handled without accounting for the inherent strategic interaction in the situation. Thus, the gaming perspective that we adopt means that we consider C2 decision-making being an activity where commanders make decisions based on their judgment regarding the other commanders' judgments and that the decision-making takes place in a C2 context.

The gaming perspective can also be derived by looking at the C2 system supporting the commander. The goal of any C2 system is to keep track of and use available information in a proper and timely manner to establish situation awareness that can be used for planning and decision-making. Situation awareness is the term coined for the state of "knowing what is going on in order to figure out what to do". To a large extent, conveying situation awareness is the utmost goal for most research and development within C2, encompassing both technical and human perspectives. The technical perspective involves the techniques and the artifacts used to establish situation awareness, typically seen as a process starting with sensory data that is successively refined into comprehensible information conveyed to the commander in form of a so-called "situation picture". The human perspective, on the other hand, takes the mental processes and the staff procedures as a starting point focusing on, e.g., how data is actually interpreted by the commander and how it should best be presented, planning models supporting the human decision process, how to share the same awareness among several commanders, etc. However, regardless of the perspective, and this is where we get back to the gaming perspective, "knowing what is going on in order to figure out what to do" inevitably includes the act of anticipating opponents' likely decisions and, in turn, what the opponent may infer regarding our decisions. That is, merely presenting a comprehensible description of the situation does not give a complete understanding of the development of a situation. Hence, a C2 system must include prediction of opponent plans and these plans are intertwined with our own plans. Equally important, of course, is to establish an appropriate mental awareness regarding how these intertwined plans depend on each other and, in turn, how to use such knowledge.

To summarize, the assumptions underpinning the gaming perspective presented in this thesis are that:

- gaming is something fundamental that characterizes all the various kinds of strategic interaction that we can think of,
- various kinds of games and game play share the same fundamental properties,
- C2 situations, e.g., disaster relief, war, etc., by necessity include gaming and, furthermore, makes it difficult to handle due to the many different, uncertain and complex factors that characterize a C2 situation.

1.2 Research Issues

This thesis investigates new means to design and improve upon computerized decision support tools in support of information and uncertainty management in C2 systems. A widely accepted fact motivating the research is that there exists a gap between existing theory and actual applications. The thesis intends to lessen this gap by approaching the problem from several directions, largely divided into applied and theoretical approaches. Applied approaches are typically centered on the end user of the envisioned tools, hence, motivating prototype development followed by field testing. Theoretical approaches typically focus on extension of readily available theory to incorporate more realistic situation modeling. Hence, validation in these two approaches consists of testing and provable correctness, respectively.

Applied research tasks approached in this thesis, i.e., targeted by the papers, include:

- development of specific algorithms for various kinds of decision support tools,
- algorithm implementation and evaluation in real settings,
- creation of generic software for gaming to be used in laboratory C2 process research, for executive training, and as a prototype decision support tool;

whilst theoretical research directions cover:

- specification of a suitable information infrastructure in support of information and uncertainty management in C2,
- adaptation of traditional inference methodology to account for multiple opposing actors,
- improvement of decision-theoretic mechanisms to account for realistic situations.

To sum up, the overall research undertaking in this thesis consists of the development of technical artifacts and procedures in order to account for strategic interaction in C2 decision-making.

1.3 Scientific Contributions

The work presented in this thesis is based on a number of publications appearing in journals and at conferences related to information fusion, decision support, command and control, operations research, microworld research, and modeling and simulation. Hence, the presented work contributes to the intersection of these areas.

The following seven papers, summarized in Chapter 5, are included in the thesis:

I. Stefan Arnborg, Henrik Artman, Joel Brynielsson, and Klas Wallenius. Information awareness in command and control: Precision, quality, utility. In *Proceedings of the Third International Conference on Information Fusion (FUSION 2000)*, pages ThB1/25–32, Paris, France, July 2000.

II. Joel Brynielsson and Rego Granlund. Assistance in decision making: Decision help and decision analysis. In *Proceedings of the Sixth International Command and Control Research and Technology Symposium (ICCRTS)*, Annapolis, Maryland, June 2001.

III. Joel Brynielsson and Klas Wallenius. Game environment for command and control operations (GECCO). In *Proceedings of the First International Workshop on Cognitive Research With Microworlds*, pages 85–95, Granada, Spain, November 2001.

IV. Joel Brynielsson, Mattias Engblom, Robert Franzén, Jonas Nordh, and Lennart Voigt. Enhanced situation awareness using random particles. In *Proceedings of the Tenth International Command and Control Research and Technology Symposium (ICCRTS)*, McLean, Virginia, June 2005.

V. Joel Brynielsson and Klas Wallenius. A toolbox for multi-attribute decision-making. Technical Report TRITA–NA–0307, Department of Numerical Analysis and Computer Science, Royal Institute of Technology, Stockholm, Sweden, December 2003.

VI. Joel Brynielsson and Stefan Arnborg. An information fusion game component. *Journal of Advances in Information Fusion*, accepted for publication.

VII. Joel Brynielsson and Stefan Arnborg. Refinements of the command and control game component. In *Proceedings of the Eighth International Conference on Information Fusion (FUSION 2005)*, Philadelphia, Pennsylvania, July 2005.

The following five papers are not included in the thesis but have had impact on it:

VIII. Joel Brynielsson. A decision–theoretic framework using rational agency. In *Proceedings of the 11th Conference on Computer-Generated Forces and Behavioral Representation*, number 02–CGF–047, pages 459–463, Orlando, Florida, May 2002.

IX. Qi Huang, Jenny Hållmats, Klas Wallenius, and Joel Brynielsson. Simulation-based decision support for command and control in joint operations. In *Proceedings of the 2003 European Simulation Interoperability Workshop*, number 03E–SIW–091, pages 591–599, Stockholm, Sweden, June 2003.

X. Joel Brynielsson and Stefan Arnborg. Bayesian games for threat prediction and situation analysis. In Per Svensson and Johan Schubert, editors, *Proceedings of the Seventh International Conference on Information Fusion (FUSION 2004)*, volume 2, pages 1125–1132, Stockholm, Sweden, June 28–July 1, 2004.

XI. Joel Brynielsson. Game-theoretic reasoning in command and control. In *Proceedings of the 15th Mini-EURO Conference: Managing Uncertainty in Decision Support Models (MUDSM 2004)*, Coimbra, Portugal, September 2004.

XII. Joel Brynielsson. Using AI and games for decision support in command and control. *Decision Support Systems*, accepted for publication.

1.4 Organization of the Thesis

The thesis is based on, and contains, seven papers. Before these papers, this introductory text gives the author's perspectives along with background information and historical notes.

The introductory text intends to put the papers in context and should be seen as the outline of the problem area that the papers target. That is, the papers contain the results while this introductory text contains the necessary background material needed for the line of reasoning. However, the paper summaries found in Chapter 5 should make the introductory text fairly free-standing.

The remainder of the chapters in the introductory text are divided as follows. Chapter 2 provides background information and initiates the line of reasoning using the work of Carl von Clausewitz as a point of reference. The main body of the work presented in the introductory text follows in Chapter 3 and Chapter 4, largely separating our efforts to investigate C2 decision-making from a practical and a theoretical perspective, respectively. Chapter 5 summarizes the included papers and discusses their contributions. Finally, Chapter 6 concludes and discusses possible avenues of approach for further work.

The appended papers are listed in a non-chronological ordering ranging approximately from applications to theory.

As will be apparent throughout the introductory text, it is the author's explicit view that ancient operations analysts as well as more recent historical remarks enlighten and enrich presentation of results that are still valid and highly topical. This should be considered merely a matter of presentation.

1.5 Acknowledgments

The author is grateful to Stefan Arnborg, Henrik Artman, Per Svensson, and Klas Wallenius for commenting on this thesis introduction and, most important, for supporting the author's PhD work continuously all the way from the time of departure up till completion.

Chapter 2

On Command and Control

This chapter contains the command and control (C2) background information that we build upon. The perspective is mostly due to warfare, ranging from the birth of modern warfare to today's information age transformation, denoting our belief that the extremes of situations, operations and analyses are captured in full by military conflict. It is also our belief that ideas applicable to warfare, in the generic perspective that we adhere to in this thesis, are applicable to a full range of other conflict arenas that in some sense can be treated as subsets of military intervention. Noteworthy, the problem of individual decision-making can be considered a one-person game against a neutral nature, although lacking some of the complexities of a true conflict situation, and can be treated as conflict using the same governing ideas (Luce and Raiffa, 1957, p. 306).

2.1 The Command and Control Dilemma

Emergency services are examples of organizations that rely on operation management from their emergency co-ordination center. Another type of organization that relies on operation management from their staff is the military, where personnel in C2 centers need to evaluate the arisen situation to give reasonable orders based on available information. The situations that occur are different from time to time and are often ill-structured. Typically, technical artifacts of various kinds are used for decision-making.

The work presented in this thesis investigates possibilities to create decision support tools that enhance C2 decision-making. For this purpose, the definition posed by Coakley (1991, p. 53), which is broad and extensively used, covers the essential properties of C2 that we are interested in:

> In general terms, C^2 is everything an executive uses in making decisions and seeing they're carried out; it includes the authority accruing from his or her appointment to a position and involves people, information, procedures, equipment, and the executive's own mind. A C^2 process

> is a series of functions which include gathering information, making decisions, and monitoring results. A C^2 system is a collection of people, procedures, and equipment which supports a C^2 process.

As indicated, C2 is a comprehensive subject that encompasses commandment of subordinates, decision-making, situation awareness, data fusion, organizational issues, and so forth – issues to which we devote the remainder of this chapter. As a consequence, there exist a number of definitions of C2 that, on the one hand, are compatible with each other but, on the other hand, are widely different depending on different focuses regarding aspects that are important for the specific target organization or for the topic the definition is intended to support. For example, in a series of work Wallenius (2002, 2004, 2005) proposes a definition of C2 encompassing the organizational task assignment structure rather than the means the commander has at his disposal for decision-making. This definition is indeed appropriate for his purpose: the design of tools for the actual execution and commandment of orders that are already decided upon. However, in our application we focus mainly on the act of coming up with a suitable decision and settle with Coakley's broad definition of C2 which nicely conveys the important aspects of our problem. That is, referring to the given definition, we will be interested in how to come up with a decision but not primarily in how it should be carried out or the human processes surrounding it. As will be discussed below, however, these two factors are somewhat intertwined and, hence, must be considered jointly.

Development, co-ordination and maintenance of progressive information system architectures for C2 are currently undertaken by military organizations and civilian emergency management organizations throughout the world. These undertakings are rightfully considered a key task to maintaining well-functioning and operational units that improve and maintain the organization's information dominance in support of its military and/or civilian objectives. Achievement depends on development and application of the latest technology for continual improvement of information systems and support of infrastructure services for operation in both high level headquarters and in the field. Hence, much of the ongoing and anticipated work is directed towards exploiting specialist skills found within civilian communications and information technology expertise. We will be interested in creating C2 decision support tools that exploit the possibilities given by these new circumstances. We believe these tools and systems should be conceived as integral parts in a C2 center which, referring to the C2 definition given by Coakley (1991), by necessity involves information gathering and equipment but also the executives and the processes they use for decision-making. Hence, it is our belief that to understand the implications of C2, the relationship between, on the one hand, the work performed by the chief and his staff and, on the other hand, the influenced real situation, must be understood.

We have discussed C2 decision-making enhancements in terms of "tools," illustrating the architecture we have in mind. We envision decision-making tools as being part of a service-oriented architecture where the commander has the possibility to use a set of tools depending on what he thinks is appropriate for the situation at hand. These decision support tools may, for example, provide the commander with means to model his problem, simulate an envisioned solution, provide means for interactive gaming in a given scenario, ask "what if?" questions regarding a particular solution, display the current situation picture in alternative ways to enhance situation awareness, share ideas with other commanders, and so forth.

2.2 Probabilities, Gaming, and Subjective Reasoning

This section tries to put the gaming perspective and its related problems into context by discussing its outlook from a historical perspective. Of particular interest for the game-theoretic discussions in Chapter 4 is the separation of different kinds of uncertainties and games into classes depending on their complexity. As we shall see, the classification of games that we discuss is generic and holds for all kinds of games. Interestingly, the Clausewitzian characterization of game complexity is strikingly close to that of modern game theory, which was not formalized until some hundred years later by von Neumann and Morgenstern (1944). That is, game theory is a formalization of the strategic interaction problems that have been discussed for a very long time.

The roots of modern conflict theory are due to 19th century military strategists, with Carl von Clausewitz (1780-1831) as its foremost representative. The Napoleonic wars made possible by mobilizing entire countries brought about new views on the importance and content of military strategy. The strategies and tactics proposed by Clausewitz were, hence, imposed by the fact that conflict had become a much more extensive and complicated undertaking than it was previously. His thoughts about military strategy were in many respects a revolt against earlier authors who, in his opinion, had concentrated solely on the problems of recruiting soldiers, using adequate armor, training, and maintenance of fighting forces. Without underestimating the importance of being prepared, Clausewitz meant that these things are as relevant to combat as the craft of the swordsmith to the art of fencing. Clausewitz introduced the more intellectually challenging task of strategic thinking by separating this subject from the earlier mentioned tactics in the following way (von Clausewitz, 1976, p. 128):

> [...] tactics teaches *the use of armed forces in the engagement*; strategy, *the use of engagements for the object of the war.*

Clausewitz's treatment of strategic thinking was further influenced by his characterization of modern warfare as something inherently complex. He indicated that war must be treated as a total phenomenon affected by a number of conflicting characteristics. He synthesized these ideas in his paradoxical and somewhat confusing "trinity," saying, in short, that war consists of the dynamic and unstable interaction between violence, chance, and rational planning. It should be noted, however, that his book contains a more wordy and vague description of the trinity that has given rise to much discussion and debate among military theorists regarding its exact interpretation, see, e.g., Villacres and Bassford (1995), and its potential meanings relative to new threats such as terrorism, see, e.g., Klinger (2006).

Working our way through the dialectic statements of Clausewitz that lead to the trinity and, later on, to the treatment of military strategic thinking, we find the following successive propositions supporting the dialogue (von Clausewitz, 1976, pp. 84–85):

18. A Second Cause Is Imperfect Knowledge of the Situation
19. Frequent Periods of Inaction Remove War Still Further from the Realm of the Absolute and Make It Even More a Matter of Assessing Probabilities
20. Therefore Only the Element of Chance is Needed To Make War a Gamble, and That Element Is Never Absent
21. Not Only Its Objective But Also Its Subjective Nature Makes War a Gamble

These four statements make up the foundation for the area of C2 decision-making, i.e., decision-making in large and realistic situations, eventually including opposing actors trying to outperform each other. We will elaborate a bit further on this issue and relate Clausewitz's 19th century view with topics within this thesis to see that the thoughts and problems are, on an abstract level, quite similar.

Clausewitz explains the strange behavior of conflict in terms of periods of inactivity, a sort of fundamental characterizing factor due to ambiguities present in 19th century warfare. Without losing Clausewitz's general idea, we may think of these inactivity periods as fundamental building blocks giving rise to the overall complexity in situations similar to war. Item 18 defines *imperfect information* as the basic cause (other than the incentive of defense being stronger than attack) to the complexity of warfare. The notion of imperfect information, as opposed to *perfect information*, means that the actors are unaware of the exact state of the world due to uncertainty regarding what actions have been undertaken. For example, the exact locations of opposing troops may not be known with certainty because only the opposing troops know what decision was actually made. It is important to distinguish imperfect information from the statement made regarding *chance* in item 20. Chance can be thought of as a dice throw and concerns solely uncertainty regarding the future, uncertainty which will be determined by nature and that will be resolved as soon as the future materializes. Imperfect information and chance form the two dimensions needed to classify ordinary leisure games into four categories. We illustrate these four classes along with some examples of popular recreational games in Table 2.1. Noteworthy, these four classes of uncertainty also provide the cornerstones in game theory with regard to computational tractability and mechanism design (Koller and Pfeffer, 1997) which will be discussed more thoroughly in Chapter 4.

	Perfect information	**Imperfect information**
No chance	Chess Go	Battleships Rock, paper, scissors
Chance	Monopoly Coin flipping	Poker Blackjack

Table 2.1: Classification of leisure games based on their dependence on two kinds of uncertainty: uncertainty regarding the chance of nature and uncertainty regarding the current world state. These two types of uncertainties are fundamental within the area of game theory.

Only one thing, captured in Clausewitz's 21st statement, remains to turn leisure games into reality. The inclusion of subjective judgments and standpoints regarding such diverse things as courage, opponent irrationality, unknown armament, opponent doctrine, etc., results in a third dimension on top of Table 2.1 denoting uncertainty regarding the actual model or game that is employed. This level of uncertainty is captured by the concept of *incomplete information*, incomplete as in not knowing what game is actually played. The notion of incompleteness must not be mistaken for the less complex notion of imperfectness. Imperfect information represents inherent uncertainty in a known model whilst incomplete information represents uncertainty regarding the model itself.

We will not try to expand Table 2.1 with a third "incomplete information dimension". Such real-world examples could be almost anything and listing them would be absurd. Instead, we use Table 2.1 as an underlying guiding principle when discussing the more realistic situations we have in mind. Still, a C2 decision situation should be thought of in terms of what it resembles the most: a game of cards; or, using Clausewitz's final phrase when discussing his 21st dialectic statement (von Clausewitz, 1976, p. 86):

> In short, absolute, so-called mathematical, factors never find a firm basis in military calculations. From the very start there is an interplay of possibilities, probabilities, good luck and bad that weaves its way throughout the length and breadth of the tapestry. In the whole range of human activities, war most closely resembles a game of cards.

An architecture based on handling incomplete information using a Bayesian game (Harsanyi, 1967–1968) will be the governing ingredient in Paper VI and VII where a C2 game component is outlined.

As it turns out, the presence of chance is fairly easy to handle both conceptually and computationally. It is simply a lottery where the expected outcome can be obtained by multiplying the probability of success with the possible gain. Imperfect information, on the other hand, is more complex both conceptually and algorith-

Type of uncertainty	Level of complexity
Chance	1
Imperfect information	2
Incomplete information	3

Table 2.2: Characterization of uncertainty in three groups, ranked in increasing order dependent on conceptual and algorithmic complexity.

mically. Not knowing the exact state of the world means making decisions that may be good in one world and bad in another. Finally, playing a real world game, i.e., not knowing the exact rules of the game, is an even harder undertaking. We summarize the various types of uncertainties we may encounter when reasoning about an uncertain situation in Table 2.2, sorted with respect to increasing complexity.

Now, let us consider the 19th statement of Clausewitz, i.e., that periods of inaction make war a matter of assessing probabilities. The two key words in the phrase are "time" and "assess". Making informed decisions is about assessing probabilities and for a decision to be evaluated there needs to be enough time to assess the probabilities. Time is no less important in the 21st century than in the 19th century; regardless whether it is a military commander or a computer tool that assesses the probabilities, enough time is still needed to be able to do the assessing. Hence, the 19th statement draws the line between making informed decisions and reacting based on skill. The amount of available time will dictate to what extent a proposed tool or routine will be used. A soldier on foot will hopefully act based on instinct when fired upon, a modern naval ship uses computerized thinking for decision support, whilst generals in a C2 center may use days for contemplation and decision-making.

Indeed, it should be noted that the work, and the propositions, of Clausewitz is rich enough to support almost anything a strategist might have in mind. This is not to be regarded as a fallacy; instead, the timeless and thorough exposition should be used for inspiration and thoughtfulness to enrich one's ideas. Although written almost 200 years ago, the work's close resemblance with today's management problems is striking. We end the Clausewitzian exposé with the following quote that we believe describes and brings together the problem area that the seemingly disparate papers, results and thoughts underpinning this thesis belong to (von Clausewitz, 1976, p. 80):

> Once the antagonists have ceased to be mere figments of a theory and become actual states and governments, when war is no longer a theoretical affair but a series of actions obeying its own peculiar laws, reality supplies the data from which we can deduce the unknown that lies ahead.
>
> From the enemy's character, from his institutions, the state of his affairs

> and his general situation, each side, using the *laws of probability*, forms an estimate of its opponent's likely course and acts accordingly.

2.3 Military Transformation

> The need to fight quickly led man to invent appropriate devices to gain advantages in combat, and these brought about great changes in the forms of fighting.
> – Carl von Clausewitz (1976, p. 127)

From time to time, technical or organizational advances spur major transitions from one military regime to another. The introduction of compulsory military service and the invention of battleships, submarines, aircrafts, radio communication, and satellites are examples of advances that have brought about such more or less profound transitions. These kinds of leaps forward in the development are typically brought about by continuous technological advances that after a while make it necessary to turn things upside-down and change the military organization fundamentally. A *revolution in military affairs* (RMA) is the term coined for such a rapid organizational and technological shift (Johnson and Libicki, 1995).

During the last decade, evolution of weapons technology, information technology, organization, and doctrine have been the motivating factors for taking a "system of systems" perspective on RMA where networked entities form the basis. The themes and levels of implementation vary from nation to nation, but whether the shift is entitled "network-enabled capability," "network centric operations," "network enabled defence," "edge organizations," or "network based defense," the foundational idea is the same: to enable widespread sharing of information by using networked capabilities. According to the vision, information sharing improves situational awareness and speed of decision-making which, in turn, enables self-synchronization resulting in improved operational effectiveness and agility (Alberts and Hayes, 2006).

Network centric warfare (NCW) is a military doctrine concept envisioned for taking the "system of systems" perspective into account by taking advantage of technical advances in information technology and telecommunications (Alberts *et al.*, 1999). The basic idea is the following (Berkowitz, 2003, p. 113):

> Network-centric warfare follows the basic idea of network-centric computing. It assumes that there is a worldwide grid of networked communications that any "platform" – ship, airplane, land vehicle, or just plain grunt – can plug into so that it can easily upload or download data. The effect is just like the Internet: what each platform happens to be is much less important than how they all work together.

From a C2 perspective, these ideas directly affect the nature of decision processes, how to allocate decisions in the organization, and the distribution of both basic data needed for decision-making and data resulting from decision-making. Hence, C2 should be seen as the core activity of the transformation and development of computerized decision support tools must account for the new possibilities and limitations that the network centric tenets offer.

Naturally, the ideas and visions of NCW are most easily implemented within the air force and among large naval ships where the platforms are already in some sense networked. Here, a few highly capable and networked platforms act in a relatively noise-free environment making automated networking capabilities tractable and fruitful. However, following the same lines of reasoning the potential gains are highest within the army where a successful NCW implementation would mean a reformation of the traditional hierarchic structure and provide opportunities for individual soldiers to act. Here, implementation is more difficult due to the noisy environment, the number of participants, and the difficulty of networking.

The basic ideas of NCW have become widely recognized around the world, but the level of implementation, definition of terms and exact content of the theme varies from country to country. A common focus for the implementation of NCW is, however, on a *service-oriented architecture*. The service-oriented concept focuses on a set of well-defined services made available on a market where actors request and offer services, more or less similar to a free market where goods and services are offered and requested based on supply and demand. The service perspective thus focuses on *what* should be provided instead of *how* it should be produced, i.e., services may very well be provided across nation borders. Put in contrast to today's organization where specific capabilities are requested from a rather static organization, this brings about significant changes that need to be implemented.

Effects-based operations (EBO) is the term coined for how one should apply NCW to accomplish overall goals. Henceforth, EBO emphasize political goals and treat military operations solely as one, out of several, possible means to reach these goals. The following broad definition has been coined (Smith, Jr., 2002, p. 151):

> Effects-based operations are coordinated sets of actions directed at shaping the behavior of friends, neutrals, and foes in peace, crisis, and war.

To treat military goals as subsidiary and focus on overall goals is not new, see, e.g., von Clausewitz (1976); Tzu (1994); but the means and opportunities to do so by using technology provided by the envisioned NCW concept are new.

The ideas behind NCW, its implementation in a service-oriented architecture, and its application in the form of effects-based operations to achieve political goals is at the heart of current military transformations. The envisioned end effect, however, is non-technical and must be evaluated on the ground of the commander's improved situational awareness which will be discussed in Section 2.4.

2.4 Awareness: Situational, Informational, Predictional

To enhance a commander's situational awareness by various means is something we will be dealing with in several respects in this thesis. Chapter 3 discusses experimentation as a means to investigate the impact that new technology has on the situational awareness while Chapter 4 discusses and proposes the actual construction of such technology.

Situation awareness is a broad term capturing almost all aspects of a person's mental awareness in a given situation. It is about being aware of what is happening around oneself and being aware of the relative importance of what is observed. It concerns awareness of situation-specific parameters, awareness of knowledge obtained from these parameters, awareness of one's possible options, awareness of possible future states and their likelihood, awareness of others' awareness, etc. Hence, people's ability to obtain situation awareness will be dependent on the kind of awareness discussed. Still, situation awareness is a fundamental and important concept in all kinds of situations that need to be controlled on the basis of mentally understanding the situation. The invention of new technology to support decision-making in, e.g., C2 centers, emphasizes the importance of the concept. Here, the invention of new technology is intended to enhance the decision-maker's mental situation awareness in order to facilitate decision-making and therefore conveying situation awareness is beset with both cognitive and technical problems that need to be considered together. On the one hand, situation awareness is the result of a mental process. On the other hand, technology intended to enhance situation awareness needs to recognize the needs of this mental process. We are interested in designing artifacts that enhance a person's, or several persons', situation awareness. To do this, we will try to infer technical requirements by studying cognitive models.

The cognitive perspective on situation awareness is a well-studied topic with Mica Endsley being a prominent representative. Many definitions of situation awareness abound, but Endsley (1988) gives a well-accepted and widely applicable one, namely that situation awareness should be described cognitively as:

> the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future.

This definition highlights three levels of situation awareness depending on the level of information refinement where (Endsley, 1995):

perception indicates basic perception of important data,

comprehension encompasses how people interpret data by combining data into knowledge, and also how people retain their state of knowledge,

projection denotes the ability to predict future events and their implication.

It is assumed that people who obtain a high level of situation awareness function in a timely and effective manner and how to achieve that, i.e., "reaching a higher level," is the primary topic on cognitive scientists' research agenda. It has been shown that people vary significantly in their ability to develop and maintain situation awareness (Endsley, 2000). Apart from differences that are situation specific, people are known to have individual cognitive abilities. Research results focusing on the cognitive aspect of situation awareness deal with human training where one learns how to develop better situational awareness either in general or for a particular purpose. Hence, the discussion in, e.g., Endsley (2000), focuses primarily on people's mental state, how people can be classified relative to some situation awareness scale, and how technical artifacts should be designed to either support people assumed to belong to a certain level of awareness or to help people go from a lower level of situation awareness to a higher level.

From a technical point of view we must be aware of the cognitive difficulties so that we design suitable technical artifacts for the whole range of commanders that we target as end users. However, an often neglected technical difficulty in the C2 context is that the technical implementation of these artifacts becomes intertwined with the cognitive aspects of the problem due to the huge amounts of data that need to be taken into account. That is, data processing in a C2 system must be made according to rules and these rules will affect the system's ability to help the commander obtain situational awareness. Taking the earlier-mentioned three cognitive levels due to Endsley (1988) as a starting point, this can be illustrated as follows:

perception of important basic data will be made *based on data that the technical solution chooses*, or has prepared in the form of aggregated data,

comprehension must be considered the ability to *use the computer to interpret and retain knowledge*, i.e., humans cannot outperform the computer when it comes to data mining,

projection should be considered a computationally and conceptually hard task where the decision-maker should be *providing basic data to be processed by inference algorithms* as outlined in, e.g., Chapter 4.

As indicated, people's development of situation awareness becomes intertwined with actual data processing. Notable correlations between the underlying technical C2 architecture and the mental state of situation awareness have, however, appeared in the area of information fusion which will be discussed further in the next section. For example, Salerno (2002) discusses the natural correlation between the level of refined data and the mental situation awareness levels posed by Endsley. Other efforts in this direction stemming from information fusion are highlighted by Bedworth and O'Brien (2000), who describe a number of existing human decision process models resulting in a new model that combines properties of the described models. Other authors, however, argue that these models, along with all other models, should serve merely as functional models that enhance common understanding and pedagogy (Llinas *et al.*, 2004). That is, there should not be a process model applied to data fusion, but rather a functional model that can be used for contemplation and understanding to enhance the development of process models needed for specific research issues or development of prototypes. In this thesis, the distinction between human decision-making versus data fusion or versus a system performing data fusion is relaxed. What is interesting is that researchers have observed the inherent cognitive problem associated with technical C2 systems based on data fusion algorithms.

2.5 Data and Information Fusion

Data fusion is a multifaceted research area dealing with aggregation and extraction of knowledge from various information sources to estimate or predict entity states. It is thought of as the core technology underlying decision support systems for crisis management, military planning and anti-terrorist applications, i.e., situations where large amounts of real-time information can be expected. Due to its versatility, data fusion is a multidisciplinary field. In a broad sense, data fusion is the process of combining data and information to gain enhanced understanding regarding the current state or the forthcoming state. All means to achieve this goal are allowed, forming an area of research where researchers from several disciplines meet.

Originally coined in 1987 by the Joint Directors of Laboratories (JDL), a U.S. DoD government committee overseeing U.S. defense technology R&D, the exact definition of data fusion is still subject to debate and continuous revisions. These revisions have caught the fact that similar underlying problems concerning data association and data combination occur in a wide range of engineering, analysis and cognitive situations. Hence, Steinberg *et al.* (1999) broadened the original sensor-centric definition of data fusion, as originally defined by White, Jr. (1987), into the following more concise definition that is the one currently in use:

> Data fusion is the process of combining data to refine state estimates and predictions.

Perhaps even more debated is the constantly revised JDL data fusion model originally outlined by the data fusion group of the JDL, see, e.g., White, Jr. (1988). This model is the most well-known and recognized method for categorizing data fusion-related processes into different levels depending on how the processes relate to the refinement of "objects," "situations," "threats," and "processes," respectively. Objects, situations, and threats can be directly thought of in terms of increased level of refinement and understanding which resemble the cognitive levels posed by Endsley (1995) described in the previous chapter. We emphasize this relationship: in C2, people's development of situation awareness is intertwined with data processing; hence, the technical area of data fusion is closely coupled to the cognitive area of developing situation awareness.

The last level in the JDL model, dealing with processes, is not really part of the hierarchical structure but is used for process refinement, i.e., it manages resources based on mission objectives and information acquired from the other levels.

A well-developed subfield within data fusion is the basic aggregation of sensor data which is called sensor fusion or multi-sensor data fusion. This area is fairly well-developed due to its relation with systems presenting sensor observations, e.g., fusion of radar plots into tracks in aviation C2 systems. Today, sensor fusion is a dependable method implemented on a regular basis by manufacturers of platform C2 systems for air and naval usage. Efficient algorithms have been developed since several decades and are well documented, see, e.g., Blackman and Popoli (1999).

Data refinement on higher JDL levels than that of sensor data fusion is denoted information fusion, typically involving various kinds of algorithms for aggregation, inference, and prediction stemming from the AI community. This thesis discusses fusion from the information fusion perspective, i.e., we deal with such things as "comprehension" and "projection" of information that, possibly, results from fused sensor reports. See, e.g., Ahlberg *et al.* (2007) for a description of a demonstrator system where a number of techniques are combined to accomplish aggregation and clustering of forces, vehicle tracking, and sensor allocation. As another example, the game component presented in Paper VI and Paper VII combines single agent uncertainty modeling techniques with game theory to accomplish higher level information fusion prediction.

Chapter 3

Simulation and Gaming

This chapter discusses simulations, games, and experiments from a point of view that brings all these activities together into the same toolbox. It is our belief that embedded simulations, in one form or another, constitute an important ingredient in future command and control (C2) system design.

Our own contribution, which will be described in the latter part of the chapter, concerns computer tools to be used for a certain class of microworld experiments, namely map based experiments that are likely to be designed as the result of research questions posed regarding the impact that new C2 technology has on the commander. Such experiments typically involve several commanders that are required to communicate to establish a common situational awareness. Microworlds highlight the importance of creating understandable models that represent the important aspects of the real situation. This model-centric view is a common important factor for the topics discussed in this thesis, be it microworlds, wargaming, probabilistic expert systems, or game-theoretic problems. In C2, it is important to make the decision-maker's mental model concrete and explicit so that it can be confronted and inspected to make it possible for the commander to apply appropriate changes to the model (Brehmer, 2000, p. 247).

3.1 Background

Simulations and games in various forms have been used for centuries by military commanders in their everyday planning activities and decision-making (Perla, 1990). A simulation is best understood as being the answer to a "what if?" question regarding an uncertain situation, i.e., by assuming values for uncertain model parameters, the simulation answers questions regarding what will happen in the model given that the assumed values turn out to be true and in the real situation if the model accurately reflects the important properties of the simulated system. Hence, it can be seen as one single play, out of possibly several, where we wish to artificially make a number of moves to see what will happen, without actually making the moves in reality. The moves are made given a model that is so hard to understand that the easiest way to get an understanding of the likely outcome is to simulate the actual move as opposed to trying to predict the outcome given the model, i.e., we use the real model in an artificial world. In this way one can try several decisions and choose the best one. Simulation is often considered as a method of last resort, but due to the complexity of the systems of interest, and of the necessary models, it often turns out to be the only way to analyze a system (Law and Kelton, 1991).

A game is similar to a simulation, but still very different because it emphasizes the strategic interaction occurring when several actors are involved in the decision-making. More precisely, the outcome for each actor, or "player," is dependent on what other participants will do – which is uncertain. Hence, the outcome cannot be determined through simulation because the variables that affect the model cannot be set in advance. That is, we can still ask "what if?" questions regarding a specific play, but the outcome will be based on our assumptions regarding the opponent's assumptions of our first assumptions and so on. This infinite reasoning loop reduces the value of asking the "what if?" question since we run a risk of being exploited by the opponent.

Games and simulations resemble each other and can be seen as application-dependent variations of the same idea. On the one hand, a simulation is just a game between nature and the decision-maker (Luce and Raiffa, 1957). This is the fundamental difference between decision theory and game theory. On the other hand, really playing a game with several reasoning actors, be it computerized or not, can be seen as a simulation that results in data and increased understanding of the simulated situation, e.g., wargaming for the purpose of preparing for battle (Perla, 1990).

The use of game-play for military decision-making is probably the oldest kind of decision support tool we can think of. Nobody really knows when or where human beings first used simulations and games for prediction of the future, but due to archaeological findings we do know that toys and games based on warlike subjects existed long before the dawn of written history. Greenberg (1981) proposes that the invention of the first wargame should be attributed to Sun Tzu, the Chinese general and military philosopher whose classic work "The Art of War" still influences and fascinates actors within all aspects of decision-making and strategic thinking (Tzu, 1994). Greenberg credits Sun Tzu for inventing a game known as "Wei Hai," probably the predecessor of the Japanese game "Go," at around the 5th century BC. Little is known about the details of the game; but similar to Go, players maneuvered armies of colored stones on a specially designed playing field and victory went to the player who managed to outflank his opponent rather than confronting him directly.

On today's research agenda is the potential use of commercially available gaming technologies for decision support (Frank and Virding, 2003). A key difference between games available commercially off the shelf and games developed especially for military use is, however, that the former are intended for the purpose of entertainment whilst the latter are used for several different purposes, e.g., analysis, training or for the support of planning and decision-making. As expected, and as we shall see, the access to the underlying model is of utmost importance in order for a game to be useful in a C2 setting.

3.2 Model Appropriateness

Both simulations and games are intimately coupled to models that capture properties of a situation, a system or some other phenomenon conceptually. To actually simulate a course of events or to play a game can be thought of as activating a model to observe, and possibly register, the course of events that takes place given the rules the model stipulates. A simulation only includes observing the development of the model's states. A game, on the other hand, requires players to interact with the model and change the model's state continuously. Still, both activities need a model depicting the situation. It follows that a well-defined and abstract model can be used for many purposes.

A model is typically a simplification of reality and, hence, two models of the same phenomenon may be different due to decisions made in the process of creating the model. Therefore, the original purpose of the model and the assumptions of the model's validity need to be taken into account by any potential users of the model. Also, there is usually a trade-off between a model's validity and its level of abstraction, i.e., on the one hand the model should resemble the modeled system and on the other hand the model should be understandable. Pure mathematical models, such as a formula saying that $distance = velocity \times time$, are the characteristic examples of the latter case. If an understandable model accurately depicts reality, we have managed to explain and understand something completely but if a simple model does not depict reality we have not gained any insights at all.

In this thesis we look upon games and simulations as activities in C2 experimentation, and upon models to be supporting these activities. Three fundamental dimensions that are underlying the logical structure in C2 experimentation can be identified, as shown in Figure 3.1, and hence need to be taken into account by the model (Alberts and Hayes, 2002, pp. 48–50):

maturity of the knowledge contribution, ranging from the discovery of new hypotheses, via the refinement of hypotheses, to demonstration of existing hypotheses,

fidelity of the experiment, ranging from wargaming, via modeling & simulation, to field studies,

complexity of issues addressed, taking into account a variety of multidimensional factors originating from the richness of the knowledge domain under study and the imagination of the experimentation team.

The overall goal of the experiment is to move toward more mature knowledge, in more realistic settings, and involving more complex issues, shown in Figure 3.1 in

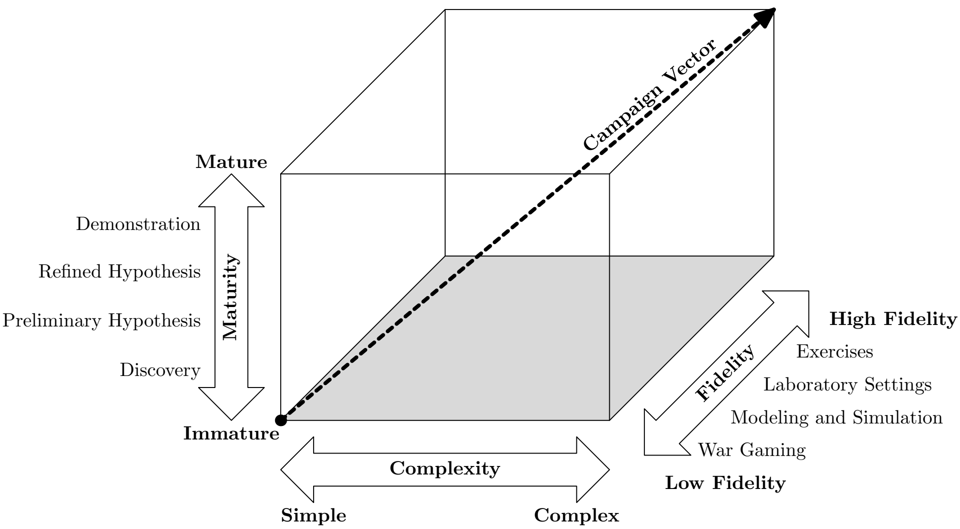

Figure 3.1: The three underlying dimensions to take into account when performing C2 experimentation: maturity, complexity and fidelity. The longterm goal is to move along the "campaign vector".

the form of a "campaign vector" that emphasizes that C2 experimentation should be seen as a campaign containing experiments on different levels along the three campaign dimensions.

In order to use models successfully in experimentation and, for that matter, any similar activity, models should adhere to the following three principles (Alberts and Hayes, 2002):

- Models must be clearly defined. The experimenter must be able to determine what is (and is not) being described in the model quickly and unambiguously.
- The contents of models must be logically consistent. The logic, algorithms, and data that describe the phenomenology of interest must be compatible. If this is not true, the "answer" the model and accompanying analysis generate could be incorrect or misleading. Seemingly simple inconsistencies can potentially have catastrophic consequences, especially in warfare.
- Models must be transparent. When models are applied and begin generating results, transparency allows team members to better interpret model behavior, identify cause and effect relationships, and obtain insights. This is especially important in C2 experimentation because emerging network-centric doctrine and enabling technologies may result in some counterintuitive outcomes that will require exploration.

3.3 Research with Microworlds

In experimental psychology, a research trend gaining importance during the last decades has been to use simplified computer-simulated worlds for experiments where several actors interact with each other. Researchers arguing in favor of such microworlds indicate that they are a means of overcoming the tension between laboratory research and field research that exist in experimental psychology (Brehmer and Dörner, 1993). In broad terms, the idea is to create a simulation that is simple enough to maintain an understandable model while at the same time being complex enough to accurately describe the studied topic.

Microworlds are not meant to provide realistic and exact simulations of physical systems. Instead, they are meant to be simplified and understandable models that preserve theoretically important criteria inherent in the modeled system. For example, researchers might wish to investigate how people handle dynamic and complex decision situations and, hence, wish to use a microworld where these properties are retained accurately. Microworlds are meant to be meaningful abstractions of a complex world and are created to make it possible to form and try hypotheses and theories and to develop these further. By that means, researchers gain increased knowledge and understanding of the studied phenomena. However, it should be pointed out, and kept in mind while experimenting, that experimental research with microworlds results in knowledge that is applicable to *theories* regarding reality, not necessarily on the reality that the microworlds represent. Microworlds do not differ from other experimental tasks in this respect and hypotheses need to be taken along the "campaign vector" in Figure 3.1 to gain generalization. If the hypothesis is neither rejected in the microworld experiment, nor in field exercises, we may say that we have a valid theory (Brehmer, 2004, p. 26). Hence, it is only through the theories one can say something about the real world, and using results from research with microworlds for direct generalization to reality would be an incorrect way to use the research method.

To be useful, a microworld should impose a recognizable task that the research subjects are intended to deal with. Typical examples include firefighting, rescue missions, counter terrorist operations, and effects based operations, i.e., tasks where the overall goal is clearly stated so that the subjects' performance can be evaluated. These situations are characterized by being *complex* in that the subjects must account for a number of different aspects, such as several different actions and several, perhaps conflicting, goals. Also, these situations are *dynamic* because the situation changes continuously and because of uncertain relations between situation-dependent variables. Moreover, these situations are *opaque* because of their characteristic black box behavior that prevent the research subjects to get access to the exact model. Complexity, dynamics and opaqueness are characteristic for situations that a microworld researcher studies (Brehmer and Dörner, 1993). These aspects characterize many realistic decision situations and microworlds have therefore been considered as a suitable research tool for studies concerning for example the development of new information technology to be used for crisis management.

As mentioned, microworld studies are often used for evaluation of the effect that new technology has on a decision-maker facing a dynamic and complex decision task, e.g., tools aiming to enhance the commander's situational awareness by various means. The microworld simulation requires the participants to form hypotheses that they try to implement when attempting to handle the situation (Brehmer, 2004). Differences in overall performance between participants using the new technology and subjects not using the new technology can then enhance our understanding regarding the usability of the proposed technology. Moreover, besides evaluating how subjects perform in the microworld, it has been shown that it is also interesting to study the actual work that is performed. For example, although performing well in the microworld, the operator may still be acting on the ground of mistaken premises (Johansson, 2005, p. 96). Therefore it is interesting not only to study the result of the work but also to study the work process itself, including possible communication between participating subjects. Studying the work process itself results in important insights regarding how an experienced operator acts in order to control a dynamic situation and, hence, results in valuable understanding regarding the domain knowledge and the skills being used for decision-making. This information can, in turn, be of great importance when developing decision support systems, i.e., experienced commanders tend to use domain knowledge and skills that are often difficult to explain for researchers unfamiliar with the domain area. Hence, it is important for a microworld researcher to be able to analyze the course of events taking place during the experiment and therefore log files and other types of experimentation monitoring tools are crucial design issues.

Although a highly specific method, microworld research has come to be recognized and used by a significant number of researchers. Examples focusing on C2 include both research and staff training. Over the years, research with microworlds has resulted in several studies and theses, see, e.g., Artman (1999); Elg (2002); Granlund (1997); Johansson (2005); Rigas (2000); Rydmark (2002), and today there exist conferences and special journal issues devoted solely to disseminating results obtained through studies performed with microworld research.

3.4 Microworld Properties

The most interesting and well-studied research task from a C2 perspective is to perform experiments focusing on dynamic decision-making. The original description of dynamic decision-making, due to Edwards (1962), describes dynamic decision tasks using the following characteristics:

1. they require a series of decisions,
2. the decisions are not independent,
3. the state changes both autonomously and as a consequence of the decision-maker's actions.

Later, Brehmer and Allard (1991) added that:

4. the decisions have to be made in real-time,

to fully capture the essence of making timely decisions.

In experimentation regarding issues in dynamic decision-making, the objective for the research subject is to *control* a dynamic and complex situation. The experimental researcher, on the other hand, wishes to investigate how humans actually perform when facing situations characterized by various degrees of dynamics and complexity. Hence, a microworld system suited for experiments regarding dynamic decision-making typically confronts the subjects with a scenario that needs to be controlled in one way or another under rather stressful conditions. As noted in, e.g., Brehmer (1992), the engineering discipline of control theory may serve as a useful metaphor for specifying the general conditions that must hold for a system to be in a state of control:

- there must be a goal (the goal condition),
- it must be possible to ascertain the state of the system (the observability condition),
- it must be possible to affect the state of the system (the action condition),
- there must be a model of the system (the model condition).

From a systems engineering perspective, we may divide these criteria into two categories where the observability condition and the action condition represent preconditions on the system whereas the goal condition and the model condition are properties of the decision-maker that the system should permit the researcher to observe.

The term "microworld" was coined by Brehmer and Dörner (1993) to denote computer experiments where subjects interact with dynamic decision problems. To fit its purpose as experimental tools where the subjects are given the task of system control, such microworlds should be designed to incorporate the three intuitive characteristics of real world dynamic decision problems that were discussed in Section 3.3: *complexity*, *dynamics* and *opaqueness*.

3.5 A Generic Perspective on Microworld Design

Research with microworlds has successfully been applied in experimental research to bridge the gap between field studies and laboratory work. However, the author of this thesis is primarily a computer scientist and, hence, his research focus has been directed towards creating suitable computer tools rather than the microworld studies themselves. The perspective of the research presented in this thesis should therefore be viewed purely as a computer scientist's interpretation of a field he is not intimately familiar with.

The key observation that has attracted our attention is that, as it seems, all microworld researchers are developing their own customized and highly specific software. Examples of such specific microworlds include:

D3Fire where the objective is to fight a large forest fire using distributed decision-making (Brehmer and Svenmarck, 1995) along with its siblings the networked FireChief (Omodei and Wearing, 1995), C3FIRE (Granlund, 1997), and the non-distributed microworld NEWFIRE (Løvborg and Brehmer, 1991),

the C.I.T.I.E.S. game (the C^3 Interactive Task for Identifying Emerging Situations), designed for studying group decision-making processes that involve dissemination and fusion of data derived from multiple sources within a metropolitan crisis control center (Wellens and Ergener, 1988),

SCUDHunt where the objective is to coordinate teams to find a certain number of hidden SCUD launchers (Perla and Loughran, 2003),

the rabbits-and-foxes task described by Jensen and Brehmer (2003) with the objective of controlling a predator-and-prey system to make it converge towards equilibrium,

Moro where the participant's task is to serve as an advisor to an African tribe who has access to a limited amount of life-supporting resources (Dörner *et al.*, 1986),

Lohhausen where the participants are required to assume the role of the mayor in a small town and rule this town for a number of years (Dörner *et al.*, 1983).

From a design perspective, these microworlds represent a wide variety of architectures, ranging from textual turn taking games to map based simulations with multiple interacting players. This is natural because a microworld is being used for an experiment that is created with a specific research question in mind. That is, the research problem comes first and the experiment is designed with the research problem in mind. Thus, it is not possible to design general-purpose microworlds and it would also be wrong to try to do so. If we did, we would be creating research problems deciding which research questions the results answer.

Due to lack of suitable software available commercially off the shelf and the researcher's wish to be in full command of the model, the tradition of quickly developing specific software for the task at hand is characteristic for the microworld research community. Moreover, the focus of the researcher, who is often not part of the software development but acting as a procurer from a software engineering point of view, is always on the forthcoming experiment rather than on the software itself, naturally resulting in software projects with large portions of old code that continuously needs to be adapted to new requirements that the researcher did not think of from the beginning. Development of microworlds on the basis of old systems may in fact endanger the experiment itself. Microworld experiments, like any experimental procedure, should be designed with a specific research problem in mind and not the other way around. Hence, thinking in terms of how to adapt an existing system is dangerous business since the research question might become of secondary interest.

Applying a software engineering perspective, in this work we have not been interested primarily in the research that can be performed with the use of microworlds, but instead the development of a generic framework that can be used for a whole class of research questions. That is, our undertaking has been to create a generic microworld framework suitable for a wide range of possible microworlds but without tempting the microworld researcher to design problems rather than questions, i.e., *first* the microworld researcher comes up with a question and *thereafter* the envisioned system *might* be of use.

The above conclusions and the results described below is the result of a two-stage process. First, an extension of a readily available microworld was performed, as described in Paper II, and, second, learning from the experiences from this work, the task of developing a generic architecture for map based microworlds was undertaken resulting in the free and well-documented Game Environment for Command and Control Operations (GECCO), described in Paper III.

3.6 Generic Software for Map Based Microworlds

Microworlds are fairly schematic compared to physically or visually realistic simulators used for operations research or for operator training. Despite their relative simplicity, microworlds are not trivial to implement. They must support multiple decision-makers, incorporate a graphical user interface, and obey a client/server architecture suited for communication of the relevant parts of the game state. However, since microworld simulations seem to use a limited set of computations, it is attractive to investigate generic microworld generators that implement the common features found in a certain class of microworld implementations. In our work we have been interested in designing a microworld framework for the wide range of map based microworlds that are likely to be needed for examining research questions regarding invention of new technology in C2 systems. Until recently, all microworld studies have been performed using custom-made software, hence, inspiring us to undertake the research task of developing a generic microworld platform. If successful, such a platform would aid many researchers in their pursuit to perform microworld studies with the opportunity to choose the exact right microworld for the research question they wish to investigate. To succeed in such a development, the platform must possess all the desirable properties of a microworld, as described previously in this chapter.

In Paper III, we describe the design and implementation of a tool, GECCO, for research in, e.g., C2, using the microworld concept. In the study we describe requirements on microworld generation along with a case study that exemplifies scenario generation in GECCO, and also shows the benefits of being able to de-

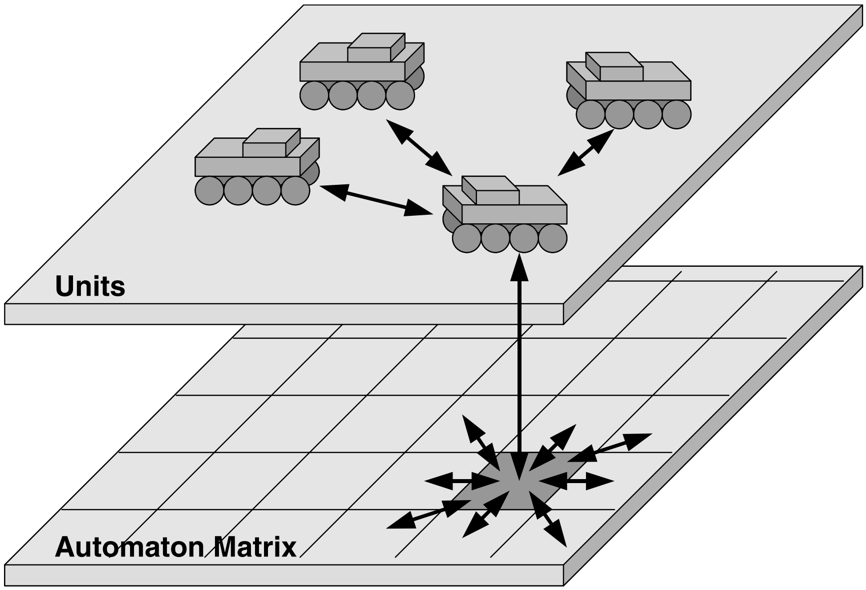

Figure 3.2: A schematic view of the layers in GECCO containing the units and the geographic automaton respectively. Arrows in the picture depict actions, implementing the interaction between the automaton and the units.

velop the code further using the concept of open source software. Games are run over a network of desktop computers with one server and a client for each participating player. The server holds the simulation and is the game engine. The clients communicate solely with the server that, in turn, tells the client what should be presented. The design is object oriented and highly abstract. The basic ingredients are a reactive automaton matrix holding the game's geographical properties, and units that move on the automaton matrix. The client/server-architecture holds the simulation and also takes care of the communication. Figure 3.2 shows a schematic architecture view of GECCO where the two layers, the automaton and the unit layer, are shown.

The GECCO framework is designed to admit rapid configuration of strategy games where decision-makers move resources on a map. Contrary to similar microworld tools, GECCO defines a game by a set of configuration files and can dynamically store all information generated during a session for off-line analysis. That is, the actual game properties, e.g., the map, the amount of players, how players interact, the properties of the geographical environment, are to a large extent defined by the researcher which makes GECCO unique in comparison to other tools used in the microworld research community. Figure 3.3 shows a potential experiment setup seen from the microworld researcher's perspective. Here, a double arrow indicates

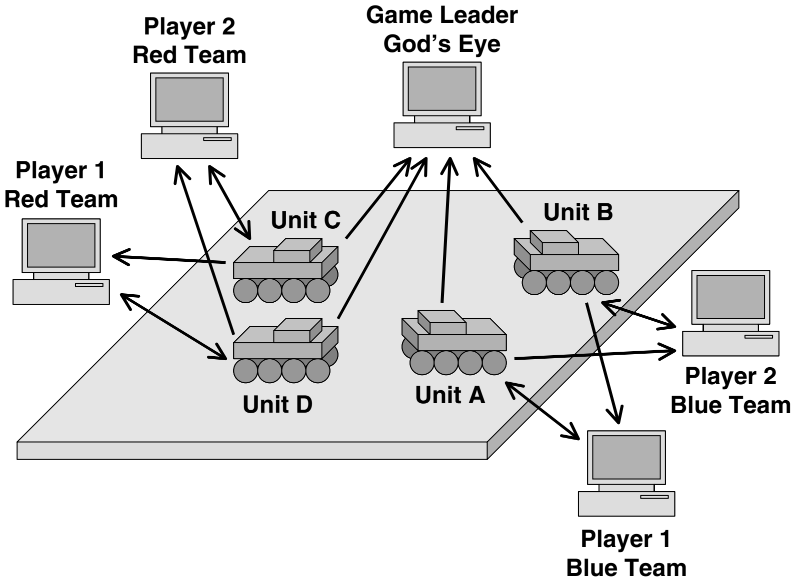

Figure 3.3: An example scenario in GECCO showing how to set up observe and command rights for players belonging to two different teams and for a game leader with a God's eye. Arrows in the picture depict information flow; a double-pointing arrow between a client and a unit means that the client both commands and observes the unit while a single arrow from a unit to a client means that the client solely observes the unit.

that a user can both control and observe a unit while a single arrow indicates that a user can only observe a unit. Hence, the picture describes a scenario with two forces where the users control one tank each and can observe all tanks in their own team.

The exact impact that GECCO has had on the microworld research community since it was developed around 2001 is not known with certainty. It is known, however, that the open source project has managed to stay alive. Every now and then users ask questions about the system and sometimes developers post updates. During recent years, a research group at the Swedish National Defence College has been using the system on a daily basis to fulfill their need for a system suited for microworld experiments. The ongoing development performed by this group on a daily basis, see, e.g., Kuylenstierna *et al.* (2004), constitutes our "proof of concept" which we illustrate using two figures. Figure 3.4 depicts a scenario map used for dynamic decision-making in a war scenario with blue and red units. Here, the researchers wish to investigate the impact of having superior sensors when facing

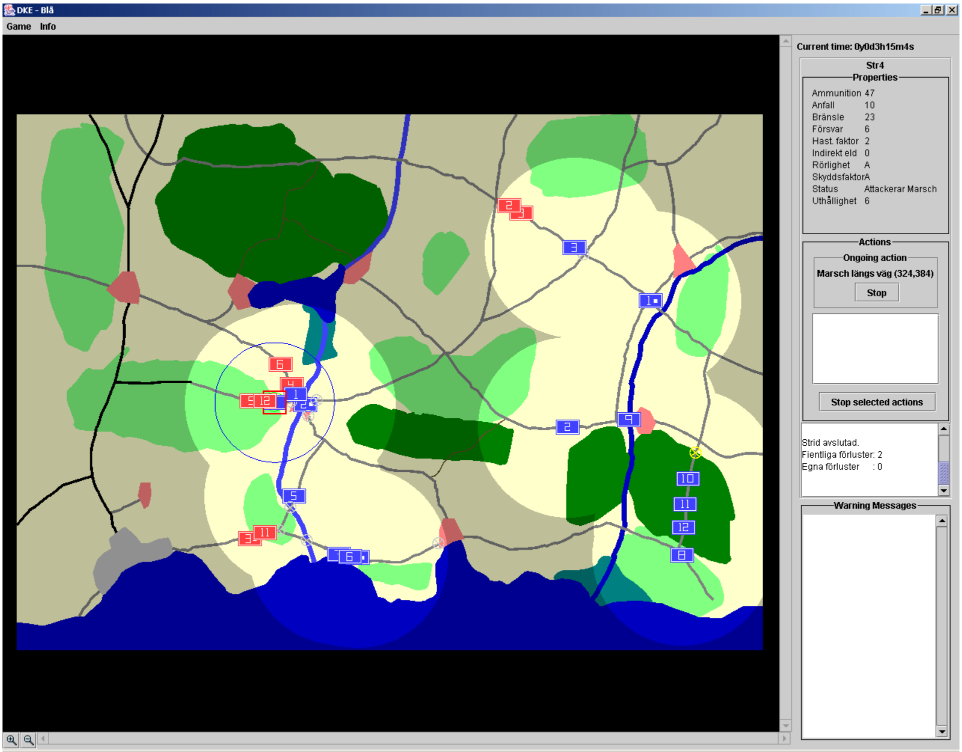

Figure 3.4: A screenshot from DKE, a wargame based on GECCO being used for experiments in dynamic decision-making at the Swedish National Defence College. The picture shows the view presented to the research subject acting as the blue commander, whose sensors can only "see" the areas surrounding his own units.

a stressful, dynamic and complex situation. The photo depicted in Figure 3.5 shows the actual experiment setup where research subjects, in this case military commanders, are trying to outperform each other in the war scenario.

The important question to ask when evaluating a microworld system is whether it is a theoretically relevant device for testing the hypotheses of interest. Hence, it is impossible to say whether GECCO, or any other piece of software, will be applicable before learning about the actual hypotheses to be tested and, moreover, it would be false to have GECCO in mind when designing an experiment. There is a dilemma coupled to this "pure" microworld research view, however. To create experiments without the slightest regard to the piece of software to be used in advance must be considered rather eccentric. After all, software projects are large and time-consuming. As it seems, microworlds have not been created from scratch

Figure 3.5: Research subjects concentrate on the experiment depicted in Figure 3.4 while being observed by a microworld researcher.

over and over again. Instead, successful results are obtained using systems under continuous development.

For the development of new C2 technology, we believe GECCO's map based architecture will be suitable for a wide range of experiments where commanders are reasoning using a situation picture. Envisioning the work performed in a C2 center as a means to convey a map based situation picture do pose some limitations on what kind of microworld studies GECCO will be suitable for, though. Nevertheless, this still captures a wide range of possible experiments that a microworld researcher needs to perform for evaluating possibilities of new C2 technology. For example, the distributed decision-making research problem (Rasmussen *et al.*, 1991) has been studied extensively using highly specific map based fire fighting microworld systems which could probably be alternatively implemented by using the GECCO framework.

Chapter 4

Decision-Theoretic Mechanisms

This chapter discusses and outlines the decision-theoretic views and tools that our work rests upon. The approach is Bayesian in that all interesting situation aspects, be it observations or hypotheses, are regarded as stochastic variables.

Readers familiar with Bayesian networks (BNs) and/or game theory probably want to skip parts of the chapter. However, we wish to highlight that the exposition differ from standard texts in at least three respects. First, it explains Bayesian inference moving straight for the BN representation by introducing the conditional probability distribution in the form an arrow at an early stage. Second, the description explains game theory using a minimum of mathematical notation. Lastly, by highlighting the possibilities and limitations that are inherent in graphical models a bridge between BNs and game-theoretic reasoning is established which, in turn, is intended to make way for the command and control (C2) game component outlined in Paper VI.

4.1 Descriptive and Prescriptive Sciences

Decision-making is studied from two perspectives: the descriptive, or explanatory, perspective; and the prescriptive, or normative, perspective (Kleindorfer *et al.*, 1993). Descriptive studies and statements attempt to describe reality, i.e., how things are in fact done. Descriptive reasoning provides the means to get retrospective understanding regarding what actually happened and why it happened. Statements resulting from prescriptive studies, on the other hand, target how things should or ought to be, i.e., valuations in the form of which actions are right and wrong and which things are good and bad. Prescriptive reasoning forms an important part of our everyday life in terms of prioritizing amongst our goals and organizing our thoughts, beliefs, emotions, and actions.

Clearly, description and prescription are closely linked in that a prescriptive analysis may benefit from feedback from a descriptive analysis, or, using a statistical viewpoint: a description may be conceived as a hypothesis that a prescription can be tested on.

4.2 Decision Theory

For analyzing a decision problem under uncertainty, the first step is to gather and structure the necessary basic data for decision-making. The following quotation, from the preface of Raiffa (1968), gives a good estimate of the necessary preparatory actions that need to be undertaken:

1. list the viable options available to you for gathering information, for experimentation, and for action;

2. list the events that may possibly occur;

3. arrange in chronological order the information you may acquire and the choices you may make as time goes on;

4. decide how well you like the consequences that result from the various courses of action open to you; and

5. judge what the chances are that any particular uncertain event will occur.

After taking these steps, meaning that the problem is systematically described and the preferences and judgments are recorded, we can start applying a strategy to propose a course of action. We base the main body of our work on expected utility which, according to Jaynes (2003), should be attributed to Daniel Bernoulli who proposed the "expectation of profit" in 1738. Notably, though, there are alternatives, as discussed in Paper V, taking for example risk into account by various more or less specific means. This chapter will, however, be based on expected utility maximization which, in combination with an elaborate model, will be our basic ingredient for handling C2 decision-making.

We now describe the traditional expected utility maximization process that we build upon. Consider every possible world, ω_j, and let every such world occur with possibility p_j. We have a number of possible strategies, i.e., the possible courses of actions, s_i, that we can perform. Further, let $u_{i,j}$ be the utility of performing strategy s_i in the world ω_j. Using this notation, the *utility matrix* is the following:

$$\begin{array}{c|cccccc}
 & \omega_1 & \omega_2 & \dots & \omega_j & \dots & \omega_n \\
\hline
s_1 & u_{1,1} & u_{1,2} & \dots & & & \\
s_2 & u_{2,1} & u_{2,2} & \dots & & & \\
\vdots & \vdots & \vdots & \ddots & & & \\
s_i & & & & u_{i,j} & & \\
\vdots & & & & & \ddots & \\
s_m & & & & & & u_{m,n}
\end{array} \tag{4.1}$$

which means that if we perform strategy s_i we achieve the following expected utility for strategy s_i:

$$E[u_i] = \sum_{j=1}^{n} p_j u_{i,j}. \tag{4.2}$$

We want to choose a strategy, s_j, to maximize our expected utility. Therefore we choose s_j so as to maximize the sum above.

Hence, traditional expected utility maximization for decision-making is described in the context of a probability distribution over the set of possible future worlds in combination with a set of strategies. Each utility value in the matrix then represents the gain given that we know which world will occur and which strategy is chosen. It should be kept in mind, however, that utility values are of interest solely because they can be compared versus each other. They do not necessarily reflect an actual measure and should be thought of as the result of a function of real consequence values, i.e., absolute values reflecting reality. That is, a utility function describes a rational decision-maker's behavior by giving a quantitative characterization of his preferences.

The combination of probability theory and utility theory constitute the general theory of decisions called *decision theory*. In decision theory, an actor is said to be *rational* if he chooses the action that maximizes his expected utility. As we will be moving on towards more elaborate structures for decision-making the underlying ideas will still be the same: a combination of probability theory and utility theory are combined into a theory that can be used for decision-making.

4.3 Probabilistic Expert Systems

This chapter explains how and why a probabilistic expert system works. We start out with ordinary probability theory as our basic building block and end up with a full-fledged inference engine. The aim has been to keep the description short and focused. Hence, mathematical rigor and interesting sidetracks have intentionally been left out in order to write an understandable and complete text that reaches its goal without losing focus. Several other books contain more detailed and comprehensive treatments of the subject, see, e.g., Cowell *et al.* (1999); Jensen (2001); Pearl (1988); Shafer (1996).

Bayesian Inference

The main property of a useful inference engine is to provide means to answer questions regarding various aspects of the world with respect to new evidence. Using ordinary probability theory, inferring something with respect to something else is captured by the concept of *conditional probabilities*. Targeting a useful inference engine instead of probability distributions for their own sake is really what distinguishes so-called Bayesian statistics from traditional practice within statistics and, hence, Bayesians consider conditional relationships being more fundamental and basic than that of unconditional joint events. In contrast, traditional practice *defines* conditional probability in terms of joint events using the standard probability axioms due to Kolmogorov (1933), i.e.,

$$P(A \mid B) = \frac{P(A, B)}{P(B)}. \tag{4.3}$$

In the Bayesian formalism, empirical knowledge will be encoded in the form of conditional probability statements, while belief in joint events, if ever needed, will be computed from those statements using Equation 4.3, i.e.,

$$P(A, B) = P(A \mid B)P(B). \tag{4.4}$$

From Equation 4.4 and the fact that $P(A, B) = P(B, A)$ we have that

$$P(A \mid B)P(B) = P(B \mid A)P(A), \tag{4.5}$$

which, by rearrangement, gives us the celebrated inversion formula named *Bayes' theorem* due to its founder Thomas Bayes (1702–1761) (Bayes, 1763):

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}. \tag{4.6}$$

The interpretation of Bayes' theorem is as follows. Suppose we are interested in the event A and that we maintain a *prior* belief $P(A)$ representing our knowledge regarding A before observing any relevant evidence. Next, suppose we observe B. Then, in accordance with Equation 4.6, we should revise our belief regarding A by multiplying our prior belief $P(A)$ by the factor $P(B \mid A)/P(B)$. The new belief regarding A given B, i.e., $P(A \mid B)$, is called the *posterior* probability. So, Bayes' theorem is all about revising belief, i.e., precisely what statistical inference is all about.

Now, the key to understand the meaning and, at the same time, the usefulness of the belief revision factor is to go back and reconsider the overall task, i.e., the act of performing inference in an uncertain world. Henceforth, let us describe our world in terms of random variables, i.e., unknown quantities that can take on one of a set of mutually exclusive and exhaustive outcomes. More exactly, let us model *all* significant aspects of the situation, be they observables or model parameters, as random variables with known distributions. For the forthcoming discussion, we reformulate Equation 4.6 using the random variables M, denoting a *model*, and D, denoting possible *data* that the model affects and that we may observe:

$$P(M \mid D) = \frac{P(D \mid M)P(M)}{P(D)}. \tag{4.7}$$

Then, for some observed value d and a possible model m, $P(d \mid m)$ by definition denotes the *likelihood* that the model m is the explanation to obtaining the observation d. Furthermore, for all possible values of M the denominator $P(d)$ will remain constant and can be treated as a normalizing constant needed to scale the right-hand side of Equation 4.7 to sum to one over all possible outcomes of M. Hence, Bayes' theorem can also be expressed by the relationship

$$P(M \mid d) \propto P(d \mid M)P(M), \tag{4.8}$$

where the expression $P(d \mid M)$, regarded as a function of m, is called the *likelihood function* for M on data d. The *causal effect* of revising belief by multiplication of the likelihood function constitutes the basic ingredient for making inferences in probabilistic expert systems. For the forthcoming discussion, also note that the likelihood function is simply a table containing conditional probabilities $P(d \mid M = m)$.

The causal effect of applying Bayes' theorem, i.e., performing multiplication of the prior distribution $P(M)$ with the likelihood $P(d \mid M)$ to obtain the posterior distribution $P(M \mid d)$, indicate a relationship where the model, M, may or may not explain the observed data, D. In Figure 4.1 (a), we display this "prior-to-posterior" process graphically. The diagram represents the structure of the joint distribution $P(M, D) = P(D \mid M)P(M)$ by decomposing it in terms of the two presumably known components $P(M)$ and $P(D \mid M)$. Hence, in the Bayesian formalism the rule $M \xrightarrow{q} D$ is interpreted as a conditional probability expression $P(D \mid M) = q$, stating that among all worlds satisfying M, those that also satisfy D constitute a fraction of size q. Following the same line of reasoning, the diagram in Figure 4.1 (b) also represents the joint distribution $P(M, D)$ using a statement saying the causal flow goes in the opposite direction, i.e., it is a statement saying that the observed "evidence" D is in reality causing the "cause" M.

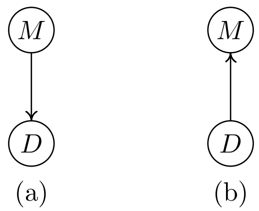

Figure 4.1: Bayesian representation of the joint distribution $P(M, D)$ depicted as (a) $P(M, D) = P(D \mid M)P(M)$ and (b) $P(M, D) = P(M \mid D)P(D)$.

Deciding what variable should constitute the model and what variable should constitute the data is potentially troublesome, but thinking in terms of "effects" and "causes" often make this come naturally. The prior probability denotes our degree of belief assigned to a model *in absence of any other information.* Hence, natural priors are things that can be based on experience and quantities known from history with clinical diagnoses serving as a typical example. For example, 30% of the adult population are known *a priori* to have a small opening between the left and the right side of the heart. This is called a patent foramen ovale (PFO) and usually has little consequence for a person's wellbeing. Clinically, however, a PFO is linked to decompression sickness (DCS). DCS incorporate the variety of *symptoms* eventually suffered by scuba divers that are exposed to a reduction in ambient pressure which, in turn, may cause dissolved inert gas to form bubbles because of the decreasing ambient pressure. Intuitively, the rate of people having a PFO, i.e., 30%, constitute a timeless fact known a priori which should constitute the model, i.e., it is the *causing* factor. DCS, on the other hand, is a good example of an *effect* that may or may not be caused by a PFO. *Given* that DCS symptoms are observed we may reason about whether they have been caused by a PFO. As indicated, Bayesian reasoning follows the problem at hand in a very natural manner, i.e., by describing the problem using natural language the Bayesian model can be deduced in a straightforward manner.

To illustrate the use of Bayes' theorem, consider a medical diagnosis problem where a doctor wishes to use the result of an X-ray to reason about whether his patient has a particular cancer or not. The result of the X-ray is known to return a correct positive result in 98% of the cases and to return a false positive result in 3% of the cases. Based on historical data, 0.1% of the population are known a priori to have the cancer. Intuitively, the cancer is the model we wish to reason about and the result of the X-ray constitutes our obtained data. The conditional probabilities for the situation can be summarized as follows:

$$\begin{aligned} P(cancer) &= 0.001, \\ P(x - ray \mid cancer) &= 0.98, \\ P(x - ray \mid \neg cancer) &= 0.03. \end{aligned}$$

Now, suppose the X-ray test returns true. Using the obtained data, the posterior probability distribution can be revised using Equation 4.8:

$$\begin{aligned} P(cancer \mid x - ray) &\propto P(x - ray \mid cancer)P(cancer) \\ &= 0.98 \times 0.001 = 0.00098, \\ P(\neg cancer \mid x - ray) &\propto P(x - ray \mid \neg cancer)P(\neg cancer) \\ &= 0.03 \times (1 - 0.001) = 0.02997. \end{aligned}$$

As mentioned earlier, normalization of the above quantities is needed to derive the final posterior probability distribution:

$$\begin{aligned} P(cancer \mid x - ray) &= \frac{0.00098}{0.00098 + 0.02997} = 0.0317, \\ P(\neg cancer \mid x - ray) &= \frac{0.02997}{0.00098 + 0.02997} = 0.9683; \end{aligned}$$

indicating that the doctor should revise his belief of his patient having cancer to be more than 30 times higher than prior to observing the positive X-ray result.

However, due to the relatively low probability for having the cancer at all, the patient is still unlikely to have cancer.

To see that normalization is valid, just notice that the data, i.e., the positive result of the X-ray, has remained constant throughout the calculation and, hence, the denominator in Equation 4.7, here $P(x - ray)$, has remained constant. The normalization process shown in this example is valid in general and brings about that the normalization constant $P(D)$ may be alternatively written using the law of total probability so that Bayes' theorem can be alternatively expressed as:

$$P(M \mid D) = \frac{P(D \mid M)P(M)}{\sum_{m \in \mathcal{M}} P(D \mid m)P(m)}, \tag{4.9}$$

where $\mathcal{M}$ denotes the set of values that M takes on. It follows that we can use Bayes' rule in two ways depending on the state of our knowledge. Either we have prior information on the data and can estimate $P(d)$ directly or, alternatively, we need to estimate $P(D \mid m)$ for all possible models m. In the example, it would have been hard to estimate $P(x - ray)$ directly, i.e., we are unlikely to have historical data regarding chest X-rays on the population as a whole. However, whether to estimate $P(d)$ or $P(d \mid \neg m)$ is inherently problem dependent. Sometimes, we have empirical data regarding the population as a whole and empirical data regarding the model that is of interest but lack data regarding the models that are not of interest. For example, the overall prior probability for developing a serious type of DCS can probably be obtained from statistics maintained by the sports diving federation, the navy and the commercial diving companies while it is less likely that these organizations maintain records indicating the probability of obtaining DCS conditioned on the non-presence of a PFO. See Bove (1998) for a Bayesian analysis of the effect that a potential PFO has on the risk of developing serious DCS, using real empirical data.

Bayesian Networks

Bayes' theorem is the basic building block underlying most modern AI systems for probabilistic inference (Cowell *et al.*, 1999, p. 14), and creating a complete inference engine is now fairly straightforward. In fact, it has already been done in Figure 4.1. Deliberately, we focused on depicting the conditional relationship pictorially using an arrow between the two random variables in Bayes' theorem. First realized by Pearl (1982), this is also the key to creating larger networks consisting of a set of random variables and their pair-wise relationships depicted by arrows. A random variable having several parents indicate there can be several factors causing a certain effect which, hence, results in that the likelihood function needs to be expanded into a *conditional probability table* (CPT) incorporating the belief of obtaining a certain effect given all possible instantiations of the parent variables. Hence, the dimension of a certain variable's CPT is equal to the number of parent-variables that affect the variable.

As an example, let us create a simplified reasoning engine for engine starter fault detection in an old car. Since the car is old and has a history of malfunctioning we have good estimates of the possible causes for breakdown. The engine starter may fail depending on either that the car battery has been discharged, that the engine starter itself is malfunctioning or both of these two factors. To model the situation, we let E, B and M be binary random variables for the events *Engine starter failure*, *Battery discharged* and *Malfunctioning engine starter*, respectively; each taking on the possible values "yes" or "no". A discharged battery causes a definitive engine starter failure whilst an engine starter malfunction in itself may or may not render the engine starter inoperable with equal probability. Also, we assume that the possible causes, i.e., a discharged battery or a malfunctioning engine starter, do not influence each other. Due to previous experience and diagnosis of engine starter failures we may assess the prior probabilities of the two possible causes. Moreover, the car headlights, which may be broken in advance, provide an additional possible symptom of a discharged battery. We let L be the random variable for the event *Light failure*. The situation is depicted graphically in Figure 4.2 with the corresponding CPTs given in Table 4.1. Note that there are four CPTs in the table because there are four random variables to reason about. Also, note that the root causes are modeled with a one-dimensional CPT containing solely the prior. It follows that a variable's prior probability distribution may be interpreted simply as a CPT conditioned on nothing at all.

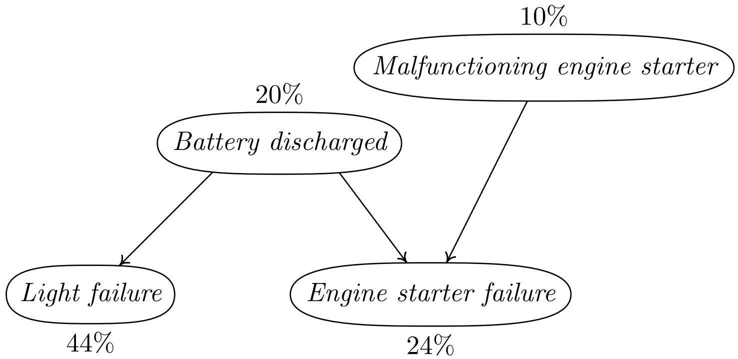

Figure 4.2: BN representing the engine starter problem before making any observations.

As said, an arrow in the BN indicates a conditional dependency. Equally important, the lack of an arrow between two variables state that these variables are conditionally independent. Henceforth, owing to the conditional independence depicted in the graph we may deduce the prior probability of an engine starter failure before making any observations by first calculating the joint distribution for the

$$
\begin{array}{llll}
B: & P(B=\text{yes})=0.2 & E: & P(E=\text{yes}\mid B=\text{no},M=\text{no})=0 \\
 & & & P(E=\text{yes}\mid B=\text{no},M=\text{yes})=0.5 \\
M: & P(M=\text{yes})=0.1 & & P(E=\text{yes}\mid B=\text{yes},M=\text{no})=1 \\
 & & & P(E=\text{yes}\mid B=\text{yes},M=\text{yes})=1 \\
L: & P(L=\text{yes}\mid B=\text{no})=0.3 & & \\
 & P(L=\text{yes}\mid B=\text{yes})=1 & &
\end{array}
$$

Table 4.1: Conditional prior probabilities for the engine starter problem before making any observations. Probabilities for "no" answers are easily calculated and, thus, omitted.

event $E =$ "yes",

$$
\begin{aligned}
P(E=\text{yes},B,M) &= P(E=\text{yes}\mid B,M)P(B,M) \\
&= P(E=\text{yes}\mid B,M)P(B)P(M),
\end{aligned}
$$

and then sum over all possible combinations of values "yes" and "no" that B and M may take on so that

$$P(E=\text{yes}) = 0\times 0.8\times 0.9+0.5\times 0.8\times 0.1+1\times 0.2\times 0.9+1\times 0.2\times 0.1=0.24.$$

Summing out irrelevant variables to obtain the probability distribution for but a few significant variables in this way is called marginalization which is an important and frequently used concept in Bayesian statistics.

In the same manner we may sum over the prior probabilities of a discharged battery to infer the prior probability of a light failure:

$$
\begin{aligned}
P(L=\text{yes}) &= P(L=\text{yes}\mid B=\text{no})P(B=\text{no}) \\
&\quad + P(L=\text{yes}\mid B=\text{yes})P(B=\text{yes}) \\
&= 0.3\times 0.8+1\times 0.2=0.44.
\end{aligned}
$$

Hence, we know all prior probabilities in the system.

Now, imagine someone turning the key to his old car realizing the engine starter is totally dead. Hence, the event $E =$ "yes" has occurred and we should update the CPT for E so that $P(E = \text{yes}) = 1$. The next step towards mending the car is to determine the cause and, hence, we need to revise our belief regarding B and M. By Bayes' theorem, the joint distribution for B and M given $E =$ "yes" is

$$P(B,M\mid E=\text{yes}) = \frac{P(E=\text{yes}\mid B,M)P(B,M)}{P(E=\text{yes})},$$

so that

$$
\begin{aligned}
P(B = \text{no}, M = \text{no} \mid E = \text{yes}) &= 0 \times 0.8 \times 0.9/0.24 = 0,\\
P(B = \text{no}, M = \text{yes} \mid E = \text{yes}) &= 0.5 \times 0.8 \times 0.1/0.24 = 0.17,\\
P(B = \text{yes}, M = \text{no} \mid E = \text{yes}) &= 1 \times 0.2 \times 0.9/0.24 = 0.75,\\
P(B = \text{yes}, M = \text{yes} \mid E = \text{yes}) &= 1 \times 0.2 \times 0.1/0.24 = 0.08.
\end{aligned}
$$

Again, the posterior probabilities for $B =$ "yes" and $M =$ "yes" can be deduced by summing the other variable out so that $P(B = \text{yes}) = 0.75 + 0.08 = 0.83$ and $P(M = \text{yes}) = 0.17 + 0.08 = 0.25$ which is reflected in the updated diagram in Figure 4.3. Also, observe that the inferred joint probability distribution has induced a dependency between B and M in that it cannot happen that there is an engine starter failure without a cause which is indeed in accordance to the conditional assumptions we entered into the system from the beginning in Table 4.1.

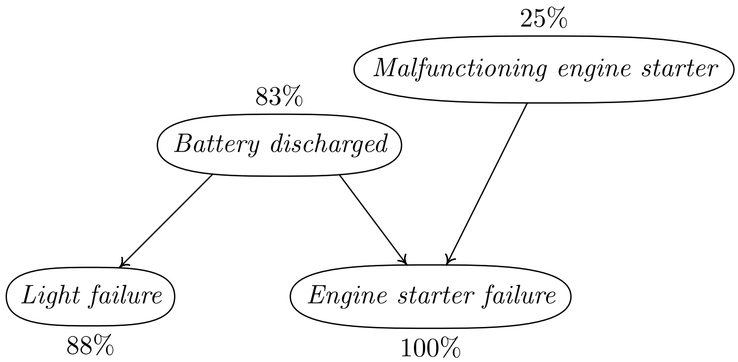

Figure 4.3: The engine starter problem after observing an engine starter failure.

To demonstrate the versatility of the BN formalism we may also observe that the probability of a light failure has gone up to

$$
\begin{aligned}
P(L = \text{yes}) &= P(L = \text{yes} \mid B = \text{no})P(B = \text{no})\\
&\quad + P(L = \text{yes} \mid B = \text{yes})P(B = \text{yes})\\
&= 0.3 \times 0.17 + 1 \times 0.83 = 0.88
\end{aligned}
$$

after observing the engine starter failure, i.e., a natural secondary cause of the increased probability of a discharged battery. The whole inferred situation arisen after observing the engine starter failure is depicted in Figure 4.3.

Now, referring to Figure 4.4, suppose we leave the key in the on position and step out of the car to observe the headlights of the car. Suppose they do not shine. Hence, the event $L =$ "yes" has occurred and we wish to shed some light on the situation using this new piece of information. In this case, the previously calculated posterior distribution for $P(B, M \mid E = \text{yes})$ becomes our prior distribution for a new application of Bayes' theorem, where $E =$ "yes" remains background evidence throughout the computation. Noting that $P(L = \text{yes} \mid B, M, E = \text{yes}) = P(L = \text{yes} \mid B)$ due to the assumed graph independence, we can derive the posterior distribution for the new situation using Bayes' theorem:

$$\begin{aligned} P(B, M \mid L = \text{yes}, E = \text{yes}) &= \frac{P(L = \text{yes} \mid B, M, E = \text{yes})P(B, M \mid E = \text{yes})}{P(L = \text{yes} \mid E = \text{yes})} \\ &= \frac{P(L = \text{yes} \mid B)P(B, M \mid E = \text{yes})}{P(L = \text{yes} \mid E = \text{yes})}. \end{aligned}$$

Hence, we obtain the following updated belief regarding B and M given the new information regarding the light failure:

$$\begin{aligned} P(B = \text{no}, M = \text{no} \mid L = \text{yes}, E = \text{yes}) &= 0.3 \times 0/0.88 = 0, \\ P(B = \text{no}, M = \text{yes} \mid L = \text{yes}, E = \text{yes}) &= 0.3 \times 0.17/0.88 = 0.06, \\ P(B = \text{yes}, M = \text{no} \mid L = \text{yes}, E = \text{yes}) &= 1 \times 0.75/0.88 = 0.85, \\ P(B = \text{yes}, M = \text{yes} \mid L = \text{yes}, E = \text{yes}) &= 1 \times 0.08/0.88 = 0.09. \end{aligned}$$

and may deduce their marginals given the new circumstances, i.e., $P(B = \text{yes} \mid E = \text{yes}, L = \text{yes}) = 0.85 + 0.09 = 0.94$ and $P(M = \text{yes} \mid E = \text{yes}, L = \text{yes}) = 0.06 + 0.09 = 0.15$. Thus, obtaining evidence in one end of the diagram affects the variables in the other end in one way or the other. In this case, the new piece of evidence supported the thesis that the battery is discharged and *explained away* the thesis that the engine starter is malfunctioning.

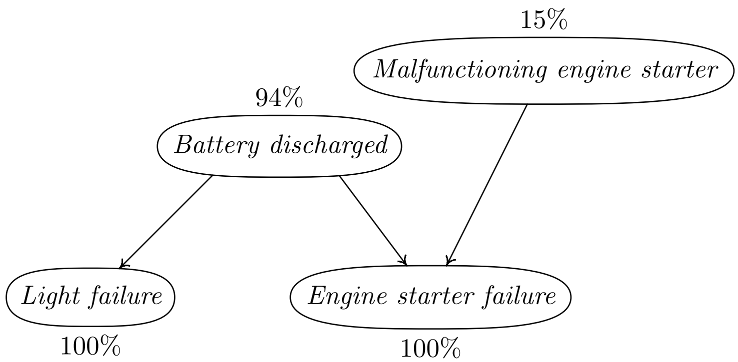

Figure 4.4: The engine starter problem after observing additional evidence stating that the car headlights do not shine.

The outlined example with the engine starter failure has argued in favor of BNs as an efficient and conceptually appealing approach for inference on the basis of further information. This is also true in general. BNs are a happy marriage between probability theory and explanatory power, lending itself naturally to the design of efficient general-purpose algorithms.

4.4 Decision-Making Under Uncertainty

As said in Section 4.2, probability theory should be combined with utility theory to form a theory to be used for decision-making. An *influence diagram* is a natural extension to a BN, incorporating decision and utility nodes in addition to chance nodes (Howard and Matheson, 1984). It represents decision problems for a single agent. Decision nodes represent points where the decision-maker has to choose a particular action. Utility nodes represent terminal nodes where the usefulness for the decision-maker is calculated.

Looking at influence diagrams the other way around, they are a means to accomplish multi-attribute decision-making. Taking the underlying idea of the BN one step further, we may exploit the utility function's conditional independencies so that the overall utility is a sum of local utilities where each local utility has as parents the attributes on which it depends. That is, letting $u = \sum_{i=1}^{n} u_i$ we create decision nodes for each u_i term and parent random variables that reflect the attributes that the function u_i depends on.

The actual evaluation of an influence diagram is straightforward; just try all possible value combinations of the decision nodes and choose the action or the actions that yield the highest utility. Efficient evaluation can be performed bottom-up by dynamic programming to obtain a sequence of maximum utility decisions.

4.5 Other Constructs for Inference and Decision-Making

Further development of the BN and its siblings is in its infancy with improvements appearing on a regular basis. Such improvements target the network structure itself, ways of learning network topology, learning network probabilities, improved inference algorithms, time-dependent models, etc. As a result, there exist a variety of constructs that resembles that of BNs and influence diagrams in one way or another. Some of the more well-known constructs are called decision networks, dynamic Bayesian networks, dynamic decision networks, and object oriented Bayesian networks (Pfeffer, 1999), to mention but a few. Apart from the network structure itself, algorithms are developed for inference and learning with many existing proposals as a result. Inference in a BN is known to be **NP**-hard for both the exact and the approximate case, as shown by Cooper (1990) and Dagum and Luby (1993) respectively, thus making it unlikely to find an efficient general purpose algorithm for probabilistic inference. Therefore, efforts are directed in several application-dependent directions resulting in various special-case, average-case, and approximate algorithms. In practice, good algorithms for inference in BNs do exist for both the exact and the approximate case, see, e.g., Russell and Norvig (2003) for a survey. Today, learning the exact network structure and its contents from data has arisen as the foremost research task (Jordan, 1999).

In this thesis we treat the various constructs for inference and decision-making as subsidiary but important. That is, we note that these constructs sooner or later boil down to the original BN, i.e., the probabilistic inference machine that we created in Section 4.3 is the key architecture that the others build upon. The inferences may be performed faster, using more efficient network structures, or using less space. In the end, however, the result of the inference remains the same and still depends on Bayes' theorem and conditional probability as its fundamental building blocks. As we shall see, this results in that the game-theoretic architecture developed in Paper VI is applicable to a number of inference engine constructs.

4.6 Graphical Models: Possibilities and Limitations

An explosion of interest in graphical models as a basis for probabilistic expert systems has resulted in a number of books dedicated to the area during the last two decades, see, e.g., the books listed in the beginning of Section 4.3. The interest is brought about by several appealing reasons making graphical modeling suitable for a number of different tasks, nicely put in the following way by Jordan (1999):

> Graphical models are a marriage between probability theory and graph theory. They provide a natural tool for dealing with two problems that occur throughout applied mathematics and engineering – uncertainty and complexity – and in particular they are playing an increasingly important role in the design and analysis of machine learning algorithms. Fundamental to the idea of a graphical model is the notion of modularity – a complex system is built by combining simpler parts. Probability theory provides the glue whereby the parts are combined, ensuring that the system as a whole is consistent, and providing ways to interface models to data. The graph theoretic side of graphical models provides both an intuitively appealing interface by which humans can model highly-interacting sets of variables as well as a data structure that lends itself naturally to the design of efficient general-purpose algorithms.

The BN is an appealing concept that can be used by a single agent for inference and for decision-making in an uncertain world. However, a BN is not suited for strategic interaction and, hence, not for the gaming perspective that we have adopted in this thesis. In the simplest and purest possible example two actors wish to reason about each other, intuitively violating the prerequisite of a BN being a directed acyclic graph (DAG). As elaborated on in Section 4.3, arrows in a BN are directly derived from Bayes' inversion formula, i.e., Equations 4.6–4.8, making the DAG property in a BN being of fundamental importance. That is, two variables in a BN cannot causally rely on each other. This should not be striking news since the BN, originating from Bayes' theorem, was created for this particular purpose: the single agent inference problem.

Due to its explanatory power, ease of computation and so forth, it would, however, be interesting to extend the BN formalism to encompass multiple reasoning agents. Successful results in this direction are found for example in the area of plan recognition where the strategic interaction is, in some cases, limited so that there does not have to be circular causal relationships, see for example the work by Johansson and Suzić (2005) or Paper IV in this thesis. However, at some point the strategic thinking becomes dominant and the BN will not be sufficient for the problem, i.e., the agent reasoning without taking the other agent's reasoning into account will be outperformed. Therefore, a BN in itself cannot take care of gaming situations, at least not in its purest possible kind. However, as shown in Paper VI, an influence diagram can indeed be used to serve as basic data for a game-theoretic architecture that, hence, maintains the nice properties of the BN. Game theory will be described in Section 4.7 as a means to bridge these two worlds together.

4.7 Game Theory

In economics, game theory is used to analyze the behavior of competing firms. It fills the conceptual gap existing between a monopoly, where a single player need not worry about what others do, and perfect competition, in which no firm is big enough for competitors to worry about. Game theory studies the intermediate case, an oligopoly, in which there are a few firms that can gain from trying to anticipate what the others will do. However, although game theory is, and will be in the foreseeable future, intimately associated with economics, there is nothing in the theory that says it should not be applied to other means of strategic interaction. On the contrary, game theory ought to be considered a science, namely *the* science about strategic interaction. Often, explanatory texts state that game theory "provides" a mathematical tool or a mathematical framework, i.e., some theory that you can choose to use as one out of several solutions for your problem. This is not entirely correct and clearly misleading. It is probably easier to understand what game theory is by explaining it the other way around: if your problem area includes "reasoning about reasoning," this *is* game theory. You cannot choose whether you have a game-theoretic problem or not; analyzing how agents interact, no matter how plain model the agents adopt, *is* game theory. Hence, game theory is about what happens when people – or nations, or firms – interact.

Game theory has had little impact on development of reasoning tools outside the field of economics, probably due to the misconception that one needs to know a large amount of fancy mathematics to *use* the concepts. However, although resting on a solid theoretical ground, actually applying the concepts is reasonably straightforward (Camerer, 2003, p. 3).

Game Theory Basics

Fundamental in game theory is the notion of a *game* which refers to any social situation involving two or more actors. A game consists of the mathematical features of a situation where a person (or a firm, or a nation, ...) must anticipate what others will do and what others will infer from the person's own actions. The specification of a game consists of the *strategies* each of several *players* have, with precise rules for the order in which players choose strategies, the information they have when they choose, and how they rate the desirability, or *utility*, of resulting outcomes. The strength of game theory is its generality and mathematical precision where the same basic ideas are used to analyze all games, no matter whether it is poker, tennis, a counter terrorist operation, or the rock-paper-scissors game.

In general, game theory poses two somewhat controversial assumptions on the game's players. A player is assumed to be *rational* in that he makes decisions consistently in pursuit of his own objectives and he is assumed to be *intelligent* in that he knows everything that the analyst knows and is able to make the same inferences as the analyst is capable of making. Hence, if we develop a theory describing the behavior of players in some game and we believe that this theory is correct, then the theory itself must assume that each player in the game understands this theory and its predictions. The assumptions of rationality and intelligence constitute the main target for criticism regarding game theory. Of course, in a real-life situation these assumptions will never be satisfied. However, theories that do not fulfill these requirements are models indicating that players can be systematically fooled, resulting in a model that is inherently fragile and will lose its validity as soon as the players learn to better understand and model the situation. As a comparison, the theoretic analysis of cryptographic protocols relies solely on the password being kept secret while the protocol itself is assumed to be known in detail by the attacker. In theory, all other assumptions would of course be naïve. Other criticism regards the simplicity of game theory, which is rarely founded in a thorough understanding of what can be accomplished by extending the actual game mechanism. Theoretically, the game mechanism can always be extended to account for (irrational) criticism regarding irrationality and, e.g., the principle of non-forgetting can be accounted for by introducing diminishing beliefs in the model. In practice, this is perhaps not as easy but the creation of more elaborate models is the key to gain increased understanding of the problem.

The exact basic constituents of a particular game depend on the characteristics of the conflict situation and the model chosen for the study. As discussed from another perspective in Section 3.2, it is important to choose a model that on the one hand is enough complex to represent the characteristics of the game and on the other hand does not hinder the analysis of the game by diminishing the fundamental issues in it. A number of models abound that takes these two extremes into account to varying degrees with the *extensive* form and the *strategic*, or *normal*, form being the most widely adopted models. The extensive form, due to Kuhn (1953) who modified the original definition given by von Neumann and Morgenstern (1944), is the most expressive and cognitively appealing form, while the strategic form and its generalization, the *Bayesian* form, are conceptually simpler forms that are suitable for general game analysis and mathematical treatment.

In game theory there is not always a "best move" that can be derived through optimization. This is part of the gaming problem where a player choosing a certain strategy with certainty is assumed to be outperformed. That is, a player that always chooses scissors in the rock-paper-scissors game would probably be outperformed. Instead, the optimum exists in the form of so-called equilibria, a form of stable solution where everyone remains as happy as possible in the given situation. In, e.g., the rock-paper-scissors game we can use game theory to infer that the best thing to do is to adhere to the strategy $1/3[\text{rock}] + 1/3[\text{paper}] + 1/3[\text{scissors}]$ which is the game's unique equilibrium solution. Here, $[x]$ denotes the lottery that always gives outcome x whereas $\alpha[x] + (1-\alpha)[y]$ denotes the lottery that gives either outcome x or outcome y with probabilities α and $(1-\alpha)$, respectively. The notation is due to Myerson (1991) and will be used throughout the chapter.

A few concepts are fundamental for calculation of game-theoretic equilibria. We give the definitions we need below.

Definition 1 (Mixed strategy) *A mixed strategy is a strategy consisting of a probability distribution corresponding to how frequently the player's pure strategies are chosen.*

Definition 2 (Best response) *A player's best response is the strategy producing the best outcome for the player, taking other players' strategies as given.*

Definition 3 (Nash equilibrium) *A strategy profile is a Nash equilibrium if each player's strategy in the profile is a best response to the other players' strategies. That is, no player has an incentive to deviate given that others' do not deviate.*

Definition 4 (Information set) *An information set for a player is a collection of the player's decision nodes satisfying that the player does not know which node in the information set has been reached when the play of the game reaches the information set.*

These definitions will be used in Section 4.8 where equilibria computation is illustrated using a concrete example. For now, observe that two players obeying the strategy $1/3[\text{rock}] + 1/3[\text{paper}] + 1/3[\text{scissors}]$ in the rock-paper-scissors game indeed play according to their best responses. However, due to symmetry all other strategies are also best responses *given* an opponent playing this strategy. However, if one player would choose to deviate to another best response this would result in that the opponent no longer played according to *his* best response which would make him change his strategy and outperform the first player. This simple line of reasoning gives some intuition for the equilibria concept as a very fragile optimal solution that on the one hand denotes a rational decision-maker's optimal solution but, on the other hand, may be inherently fragile. The fragility necessarily poses concerns regarding the actual meaning of optimality in a gaming situation.

The rock-paper-scissors example also indicated another more concrete difficulty inherent in game theory: we *proposed* the euilibrium solution and *confirmed* its validity, but we did not say anything about how we managed to *compute* the proposed equilibrium solution. Guessing the equilibrium in the rock-paper-scissors game, i.e., a symmetric zero-sum game, might be straightforward, but in general the opposite holds. Often, there are multiple equilibria that are hard to find and to classify with regard to desirability.

So, is game theory a descriptive or predictive theory as discussed in Section 4.1? Can game theory be used to predict what people do or to give people advice? The theorist's answer is that game theory is simply analytical and gives answers to mathematical questions regarding what rational players will do given that they follow a given model precisely. If people do not play according to the model, their behavior has not proved the mathematicians wrong, any more than finding that an ice cream salesman giving you the wrong change should disprove arithmetic.

In practice, however, analytical game-theoretic tools are indeed used to predict, and also to explain and prescribe (Camerer, 2003). In the AI community, game theory has emerged as the number one formalism for the study of both non-cooperative and cooperative interaction in multi-agent systems. Typically, the classical work within game theory provides rich mathematical foundations and equilibrium concepts that are used for further development into computational and representational formalisms that scale up and provide means for prediction as desired.

4.8 An Example Scenario

We illustrate the use of game theory with a small example of a conflict scenario. The scenario has been constructed with help from a military domain expert to be useful for specific research rather than being an example of a realistic C2 decision situation. A scenario map depicting the relevant units for the decision problem can be seen in Figure 4.5 along with the terminology we use when referring to the map. For simplicity, and as illustrated in Figure 4.5, throughout the scenario description we refer to friendly and hostile forces as “blue” and “red” forces, respectively.

Northern Sweden 2020

Tension has grown gradually in the Baltic Sea during the last years. As a consequence, the Swedish armed forces have been provided resources to maintain units that are on continuous alert. At the out-break of the invasion, a number of events happen at the same time, some that are immediately considered as threats and some that may or may not be threats.

At 01:00 an enemy force disembarks 50 km north of the city of Härnösand. Reports from civilians make it likely that the force consists of two heavily armed tank companies and one mechanized infantry fighting company. At about the same time, the coast guard reports that a large leakage of oil has been discovered in the region and that several unidentified cargo ships have been sighted heading towards Härnösand. Our own troops, about the size of a battalion, reside in the proximity of

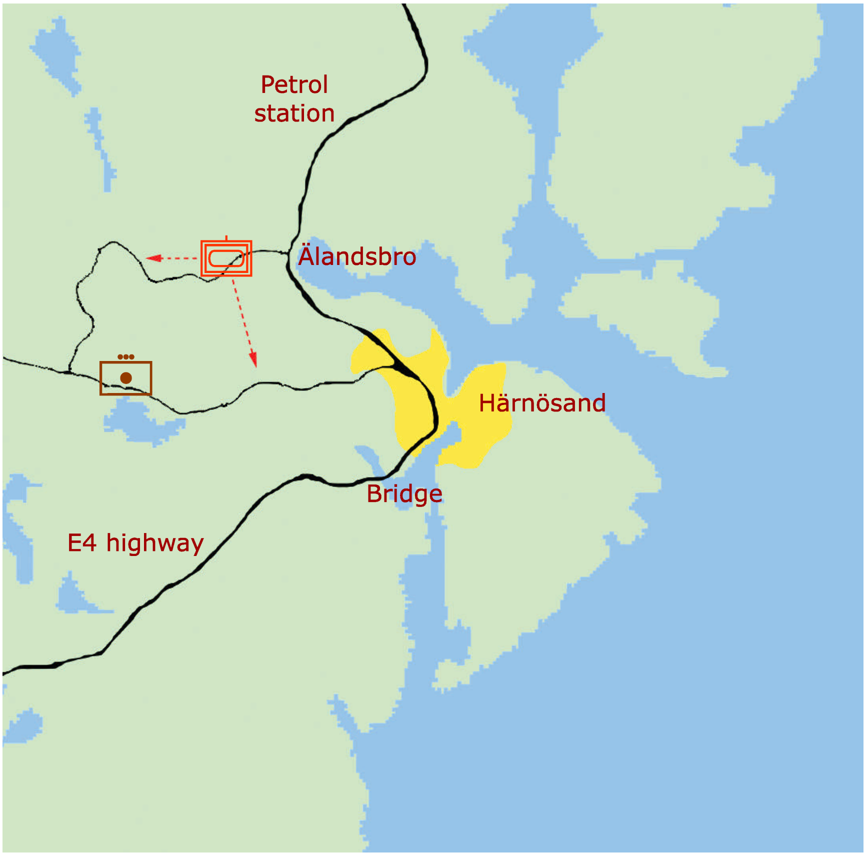

Figure 4.5: The map describes the blue battalion commander's view of the situation at the time for decision-making. The dotted arrows depict the red tank formation's choice to continue along the road versus its (potential) choice to travel through the terrain.

Härnösand. The forces mainly consists of one tank company protecting the highway E4 approaching Härnösand from the southwest and one artillery formation located along the small road directly to the west of the city. Apart from these units, a staff company together with various resources for surveillance and reconnaissance are located in the center of Härnösand. Moreover, home guard patrols of varying size and equipment capabilities may be deployed locally on various places throughout the region.

The (blue) commanders in Härnösand estimate, almost instantly, that the overall (red) enemy goal is to establish a bridgehead by gaining control over the harbor in Härnösand. This is a natural assumption as the Härnösand harbor is the only port in the area that allows for big vessels to approach in order to set off a large-scale invasion. The battalion's unmanned aerial vehicle (UAV) group is ordered to perform reconnaissance with focus on the main roads leading to Härnösand, i.e., the possible avenues of approach for other possible enemy units.

At 01:20 one of the two available UAVs spot a tank company southwest of Härnösand heading northeast on the highway E4. The battalion commander makes the assumption that the main goal for the tank company is to secure the bridge along the E4 to make it possible for more units to approach the city from the south at a later point of time.

At 01:30 a home guard patrol reports that enemy tanks have taken position at a petrol station located north of the city. At the same time the earlier-mentioned bridge along the southwestern part of the E4 is being blown up by another home guard patrol.

At 01:32 the artillery unit located to the west of Härnösand directs fire towards the petrol station using coordinates that have been supplied by the home guard patrol. The home guard patrol later reports that a big fuel explosion has made the petrol station blow up as well as neutralizing several enemy tanks that were located within the petrol station at the time of the explosion. The tanks that are not damaged, about the size of a platoon, continue south.

At 01:45 reconnaissance personnel at the crossroad in Älandsbro, see Figure 4.5, report that the remainder of the blown-up enemy tank company from the petrol station heads west while the intact enemy tank company continues south towards Härnösand.

The Need for Reasoning

At this point the blue battalion commander in Härnösand faces his first real problem where he needs to reason about the situation. Until now the commanders have been faced with several tasks that require the use of their units, C2 system, etc., but these tasks have all been straightforward with few real choices for the commanders to reason about.

The 01:45 report from Älandsbro, however, makes it apparent that the intact enemy tank company heading for Härnösand is approaching the city to, at one time or another, enter the city to seize control over it. It is also apparent that the enemy tanks heading west are left behind to take care of the blue artillery unit that, in the red commanders' view, needs to be neutralized. A fundamental principle in war, stated in the doctrine of most modern armies, is to never leave enemy forces behind as one advances.

The game-theoretic decision situation that has arisen concerns the red platoon heading west and is due to the incident at the petrol station that may indicate that the red tanks were in desperate need of fuel. There may, however, be other reasons to why one stops by at a petrol station. Moreover, even if there was a problem with the fuel, the tanks were maybe able to refuel before the petrol station was hit by artillery fire. The question regarding fuel becomes important when about to give orders to the blue artillery unit regarding his course of action. If a tank is low on fuel it must use roads because of the many times higher fuel consumption required when moving through terrain.

The blue battalion commander has two possible courses of actions to propose to his artillery unit. Either he can prepare for a certain battle by *digging trenches* or he can choose to *withdraw* towards Härnösand using the road leading to the east. Digging trenches will compensate for the heavy enemy fire power and yield an even, but certain, battle. Withdrawal, on the other hand, is more risky. If the red tanks use the road and the blue artillery unit manages to get away, this will yield a small profit for the blue side who can gain from the information and people that are saved. However, if the blue artillery unit chooses to withdraw and the red tank platoon travels through the terrain and manages to intercept, the blue artillery unit will be annihilated and their weapons might be used for the forthcoming invasion of Härnösand.

The red commander is neither uncertain about the status of himself (out of fuel or not) nor the status of the blue unit. The eventual choice to either travel using the *road* or to go through the *terrain* has got to do with that he does not know whether the blue artillery unit chooses to dig trenches or to withdraw. Also, even if the tanks are not out of fuel it should not be taken for granted that the best choice is to use the terrain to accomplish an interception, as the red tank platoon will then use a large amount of his fuel supplies which may prove fatal later on. Several historical examples of the out of fuel scenario exist, with the sinking of the German battle-ship Bismarck during WW2 being a famous example (Durschmied, 1999).

4.9 Solving the Example Scenario

The Härnösand scenario possesses several levels of uncertainty which makes the situation ideal for game-theoretic reasoning. First of all, the blue player is uncertain regarding what game is actually being played, i.e., whether the red player is out of fuel or not. Modeling this prior information requires the use of a Bayesian game. The Bayesian property is often modeled using a historical chance node as a root node. This node differs from an ordinary chance node in that the outcome of this node has already occurred and is known to a subset of the players when the game model is formulated and analyzed. In our example there are only two edges going out from the root node. One of these edges corresponds to the extensive form game in Figure 4.6 (a) that models the situation when the red player has got enough fuel. The other edge corresponds to the extensive form game in Figure 4.6 (b) that models the out of fuel situation.

The uncertainty regarding the other player's decision is modeled via the use of information sets, i.e., the red player will not know in advance whether the blue player has chosen to dig trenches (D) or to withdraw (W) and therefore is uncertain

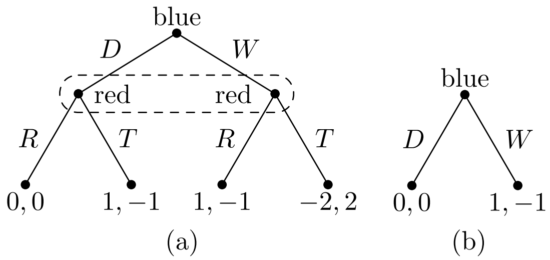

Figure 4.6: The resulting extensive game when the red force (a) has enough fuel and (b) when the red force is (almost) out of fuel.

about whether to try to intercept (T) or not (R), as his reward from this differs depending on the blue player's actions. The final expected payoffs depend on the opponent's beliefs and on our beliefs. The underlying leaf node payoffs represent the payoffs for all possible outcomes of the game. In the model outlined in Paper VI, these payoffs are obtained from the influence diagram that represents our current situational awareness.

We let $\alpha \in (0, 1)$ denote the blue player's belief of the red player having enough fuel. Solving the game using the technique described by Harsanyi (1967–1968) involves introducing a historical chance node, a "move of nature," that determines the red player's type, hence transforming the blue player's incomplete information about the red player into imperfect information. The Bayesian equilibrium of the game is then precisely the Nash equilibrium of this imperfect information game. The Harsanyi transformation of the Bayesian game is depicted in Figure 4.7 on extensive form. Note that several decision nodes share the same label representing the uncertainties regarding players' types and choices. The normal way of solving such a game is to look at the strategic representation, as seen in Table 4.2.

To solve the game, we first look for equilibria in pure strategies. $([D], [T])$ is not an equilibrium because at this outcome there is an incentive for the red player to change her action to R. The strategy profile $([D], [R])$, however, is not a stable point in the game because here the blue player can benefit from performing W instead of D. Once again, though, $([W], [R])$ cannot be an equilibrium of the game since the red player can get $3\alpha - 1$ instead of only -1 by performing T instead of R. At $([W], [T])$, a requirement for the blue player to be unwilling to change from W to D is that $1 - 3\alpha \geq \alpha$, which is true for $\alpha \leq 1/4$. Hence, for $\alpha \leq 1/4$ the pure strategy profile $([W], [T])$ is the unique game equilibrium, and for $\alpha > 1/4$ there are no equilibria in pure strategies and we have to look for equilibria in mixed strategies.

For $\alpha > 1/4$, we let $q[D] + (1 - q)[W]$ and $s[R] + (1 - s)[T]$ denote the equilibrium strategies for the blue and the red player respectively, where q denotes the probability that the blue player digs trenches and s the probability that the red

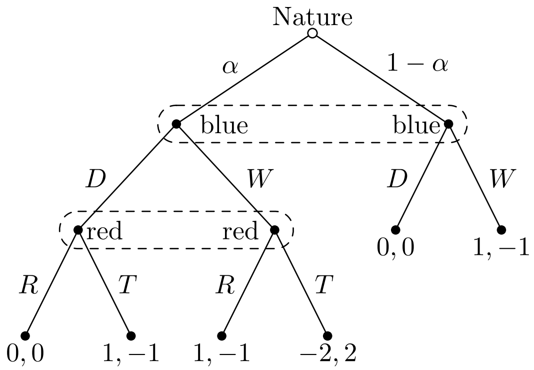

Figure 4.7: The Harsanyi transformation of the Bayesian game represented by the type profiles depicted in Figure 4.6.

	red	
blue	R	T
D	$0,0$	$\alpha,-\alpha$
W	$1,-1$	$1-3\alpha, 3\alpha-1$

Table 4.2: The strategic form of the game in Figure 4.7.

player decides to travel on roads. A requirement for an equilibrium for the blue player is that his expected payoff is the same for both D and W, i.e.,

$$s \times 0 + (1-s) \times \alpha = s \times 1 + (1-s) \times (1-3\alpha) \Rightarrow s = \frac{4\alpha - 1}{4\alpha}.$$

Similarly, to make the red player willing to randomize between R and T, R and T must give her the same expected utility against $q[D] + (1-q)[W]$ so that

$$q \times 0 + (1-q) \times (-1) = q \times (-\alpha) + (1-q) \times (3\alpha - 1) \Rightarrow q = \frac{3}{4}.$$

Since this determines the value of both q and s uniquely, there is exactly one equilibrium point in the game for all values of α, which is also a property of all constant-sum two-player games (Chvátal, 1983; von Neumann, 1928).

We can now use the equilibrium strategy of the imperfect information game to derive the Bayesian equilibrium of the Bayesian game. A Bayesian equilibrium specifies a randomized strategy profile containing one strategy $\sigma_i(\cdot \mid t_i)$ for all combinations of players and types. Hence, the unique Bayesian equilibrium for the whole game is

$$\begin{aligned}
\sigma_{\text{blue}}(\cdot \mid \text{blue}) &= q[D] + (1-q)[W] = 3/4[D] + 1/4[W], \\
\sigma_{\text{red}}(\cdot \mid \text{red.enough fuel}) &= s[R] + (1-s)[T] = (4\alpha-1)/4\alpha[R] + 1/4\alpha[T], \\
\sigma_{\text{red}}(\cdot \mid \text{red.out of fuel}) &= [R],
\end{aligned}$$

for $\alpha > 1/4$ and

$$\begin{aligned}
\sigma_{\text{blue}}(\cdot \mid \text{blue}) &= [W], \\
\sigma_{\text{red}}(\cdot \mid \text{red.enough fuel}) &= [T], \\
\sigma_{\text{red}}(\cdot \mid \text{red.out of fuel}) &= [R],
\end{aligned}$$

for $\alpha \leq 1/4$. The solution graph is depicted in Figure 4.8, showing how the equilibrium probabilities varies for different values of α.

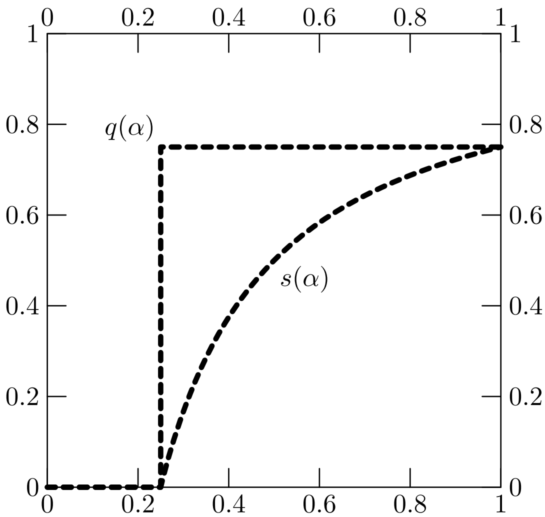

Figure 4.8: The graph shows how two players' equilibria, with probabilities indicated by $q(\alpha)$ and $s(\alpha)$ respectively, vary depending on their prior beliefs about the other player's private information α. The players differ in that the first player is speculating about what the second player knows, whilst the second player is speculating about the first player's speculation.

From the red player's perspective, the intuition is that she should try to intercept with a greater probability when she is more likely to achieve a surprise effect, i.e., when the red player thinks the blue player thinks she is out of fuel she should, if possible, try to intercept with a greater probability. For $\alpha \leq 1/4$ the surprise effect is large enough to make the red player always try to intercept if her fuel permits.

From the blue player's perspective, his belief of the red player being able to intercept, i.e., α, and the probability that the red player actually tries to intercept, i.e., $s(\alpha)$, outweighs each other so that he tries to dig trenches or withdraw according to constant probabilities q and $1-q$ for $1/4 < \alpha \leq 1$. For $\alpha \leq 1/4$, i.e., when he does not expect the red player to be able to intercept, the expected utility of a withdrawal becomes so large so that he always tries to withdraw.

It is important to consider the Harsanyi transformed game matrix in Table 4.2 solely as an intermediary result which is used to solve the game. When analyzing the game, its origins must be taken into account. If one tries to interpret the Harsanyi transformed game matrix without considering the game's origins, it probably feels strange that the game matrix contains utility values that are functions of α, i.e., a variable describing the blue player's belief of something the red player already knows with certainty. Instead, considering the origins of α, one should interpret the Harsanyi transformed matrix as the model that gives rise to the rational course of action given a subjective judgment of α. Hence, for the blue player α constitutes his own estimate regarding the red player's fuel situation. For the red player, α constitutes her estimate of the blue player's estimate. Therefore, despite that the true value of α is already known by the red player, α will still affect the blue player's rational course of action and, hence, also the red player's rational course of action.

4.10 Solution Interpretation

Nash equilibria, in the form of mixed strategies, as a solution to decision problems require a moment of thought. On the one hand, it is easy to argue that the equilibrium strategy is theoretically sensible. After all, the notion of Nash equilibria, building on the concept of rationality, defines precisely this. By using the idea of Bayesian games we are able to create alternative models regarding agents that are in some way "irrational". Thus, by using Bayesian games we can counterattack objections on the existing model by extending the model with a new sub model that models the objection in question. Of course, this also requires assigning a prior probability to the new sub model and re-evaluating the prior probabilities for the existing sub models, which makes sense if someone comes up with an objection (which is interpreted as a new model that we have not thought of before). If the objection is independent of the existing models, normalization is the natural way to re-assign probabilities. Otherwise it is natural to let the prior probability of the new model be represented by a reduction of prior probabilities of the model or the models that it depends on. In most cases we believe that it is appropriate to have a separate model for the "uncertain case" that takes care of whatever we have not thought of. In that case the new model, provided it is independent of other existing models, typically reduces our overall uncertainty regarding the situation and thus causes a reduction of prior probability for the earlier mentioned "uncertain case". Models that takes care of the rest, i.e., that represent options or possibilities that we are not yet aware of, are often found in proposed architectures for multi-agent modeling, see for example Gmytrasiewicz and Durfee (2000) where irrational behavior as well as lack of information is modeled in so called "no information models".

On the other hand, although representing the theoretically rational course of action, the Nash equilibrium poses several concerns regarding its interpretation. Looking at the Härnösand scenario, it is interesting to see how q and s vary depending on α which is depicted in the diagram in Figure 4.8, i.e., how the solution to our decision problem varies significantly depending on our subjective beliefs regarding the out of fuel situation. How do we convince a commander that he should decide what to do by throwing a die that varies depending on $q(\alpha)$? As an example, consider the situation when we do not know anything and assign equal probabilities to the two models (fuel or out of fuel). Then the blue player should dig trenches with probability $q(1/2) = 3/4$ and the red player should choose to travel by roads with probability $s(1/2) = 1/2$ although his fuel supply would allow for an interception. The conclusion regarding the Härnösand scenario is that a simple problem yields a solution that is difficult to understand intuitively. Unfortunately, this is quite typical, see for example Paper VI for another example, and we need to address the question of how to use the solution in a sensible way. To actually throw the die is part of the solution and if this is not performed the commander is not rational and, hence, will be outperformed by a rational opponent that is capable of modeling this behavior. Maybe it is easier to accept the opponent's randomized strategy as a prediction. Then the optimality of one's own randomized strategy is fairly easy to establish.

A modern interpretation of mixed strategy Nash equilibria stemming from behavioral game theory is that players need not actually randomize, as long as other players cannot guess what they will do (Camerer, 2003). Rather than considering the opponent's strategy being the throw of a die, one considers "an equilibrium in beliefs". The underlying idea is to consider the opponent being part of a population of decision-makers who choose their strategies according to a frequency that corresponds to the mixed strategy equilibrium. Hence, the opponent's choice of strategy will coincide with his equilibrium strategy on average, making the players indifferent about which strategy they play. While the human perspective that characterizes behavioral game theory is opposite to that of the C2 decision-making perspective, this modern interpretation remains a tempting way to reason about an equilibrium solution.

4.11 Computational Issues

Although the example in Section 4.9 was fairly easy to solve, it should be noted that this is often not the case. Solution methods for game-theoretic problems are, in most cases, intractable for the generic case. Two-player zero-sum games form the exception. Since the two players' payoff matrices, A and $-A$ respectively, are identical apart from the sign, the problem of calculating optimal mixed strategies x and y for the row and column player respectively can be solved via a polynomial time linear programming algorithm where the row player aims to maximize and the column player aims to minimize the sum $x^T Ay$. Apart from being computationally easy to find, the celebrated Minimax Theorem (Chvátal, 1983; von Neumann, 1928) states that the resulting equilibrium point is unique.

The most well-known solution method for general-sum two-player games, the Lemke-Howson algorithm (Lemke and Howson, Jr., 1964; von Stengel, 2002), solves a linear complementarity problem (Cottle *et al.*, 1992). The computational complexity for finding one equilibrium is still unclear. We know, according to Nash's theorem (Nash, 1951), that at least one equilibrium in mixed strategies exists but it is problematic to construct one. Lemke-Howson exhibits exponential worst case running time for some, even zero-sum, games. However, this does not seem to be the typical case. Interior point methods that are provably polynomial are not known for linear complementarity problems arising from games (von Stengel, 2002). Methods amounting to examining all equilibria, such as finding an equilibrium with maximum payoff, have unfortunately been proven **NP**-hard (Gilboa and Zemel, 1989).

The perhaps most frustrating, yet challenging, observation for the applied game-theorist is that most research and development seems to be directed towards finding *one* equilibrium instead of finding them all. Moreover, if one manages to find all equilibria the problem remains to rank these equilibria versus each other. Ranking strategies exist, but they are ambiguous and do not provide a definite answer for all situations. Ranking methods coupled with solution methods as a means to reduce the optimization problem in order to get hold of the right equilibrium are yet to be developed.

Chapter 5

Summary of Included Papers

In this chapter the included papers are summarized and their respective contributions are highlighted.

5.1 Paper I: Information Awareness in Command and Control: Precision, Quality, Utility

The paper deals with the value of information and coins the term "information awareness" as a means of being aware of the uncertainties inherent in the awareness of a situation and, hence, encompasses that there has to be an understanding of the usefulness of information and the possibilities to achieve better information.

To maximize the total benefit of the information resources, a means of measuring the usefulness of the information is required. Decision-makers should always be aware of to what extent they can trust the information, and what information they have, compared to what information they need in their current assignments. They should also, somehow, be aware of how they could benefit by using more of the information resources. Thus, we state that information awareness must be included in the concept of situation awareness.

Also, three measures are defined that, if presented along with the information, will give a larger degree of information awareness to the users. *Precision* denotes measures of the "correctness" of data, *quality* denotes fitness for purpose, and *utility* denotes the expected benefit for using the piece of information.

The analysis made in the paper is Bayesian, the view that all kinds of uncertainties can be described as probabilities. Interestingly from the author's PhD student perspective, it is noted in this first paper that "the opponents' intentions can not yet be taken care of in an adequate way" which is a game-theoretic question dealt with extensively during his last years as a graduate student, see, e.g., Paper VI and VII.

5.2 Paper II: Assistance in Decision Making: Decision Help and Decision Analysis

The paper outlines two main directions for the development of decision support tools, *decision help* and *decision analysis*, and defines their exact meaning to be used within the scope of the paper. Decision help, on the one hand, is envisioned as a tool that helps the commander look at the whole picture in order to suggest one or several courses of action. Decision analysis, on the other hand, is defined as a tool that gives the commander feedback and suggestions of improvements regarding a specific course of action that he has already decided upon. Both decision help and decision analysis are thought of as being implemented in computerized tools connected to command and control (C2) systems.

Referring to research performed at the Swedish National Defence College, it is pointed out that decision help should not be given a status of great importance. Instead, it is important that decision-makers involve themselves actively in the decision process and decide what to do without regard to decision proposals obtained from automatic tools. Humans are good at finding general patterns, but are not very good at analyzing details. Hence, the paper emphasizes that tools for criticism and improvement of decisions that are already decided upon should be considered useful.

Finally, the applied part of the paper describes a generic rule based algorithm for decision analysis in maps which has been implemented in Java. The algorithm was integrated and tested within a map based microworld system where the objective is to control large forest fires. The actual analysis performed concerns fire-fighting unit leaders' proposed movements of fire extinguishing units with regard to the conveyed situation picture.

A lesson learned in this project, not explicitly stated in the paper, was that the microworld integration of the developed tool would not have been possible were it not for assistance from the maintainer of the microworld himself. This insight along with some surveying revealed that microworld researchers are stuck with highly specific microworlds that are hard to maintain and improve, hence, stimulating the creation of a generic and maintainable framework intended to be suitable for a wide range of map based microworlds, which is the topic of Paper III.

5.3 Paper III: Game Environment for Command and Control Operations (GECCO)

The Game Environment for Command and Control Operations (GECCO) is a generic framework for map based microworld design and gaming. It is built in Java and is very generic, meaning it is easy to add different types of scenarios without changing the game and also that the code itself is easy to improve to meet the researcher's demands. An underlying assumption in the paper is that commercial strategy war games available off the shelf are not well suited for research purposes, because they cannot be changed since the source code is not freely available. GECCO is free software released under the GNU General Public License (GPL).

GECCO uses a client/server solution. The server keeps track of the simulation and any number of clients can connect to the game. The client itself does not know anything about the currently playing game; it is just a graphical user interface. The architecture is divided into two layers: one layer where units can move around and another layer representing the background. The unit layer is vector based, so the units can move in any direction, while the background layer consists of an automaton matrix. Each cell is a square block containing a color, representing forest, water, mountains, etc. The automaton feature makes it easy to simulate, for instance, a forest fire that widen to adjacent squares depending on applicable rules.

In GECCO, the microworld research community is equipped with a tool that has not been available before, facilitating research on both behavioral and technical questions. Its scope, although not mentioned in the paper, is not limited to C2 research. The main property that makes GECCO a suitable tool for C2 research as well as other research is that it is able to represent a large genre of spatially oriented games. Moreover, GECCO is: constructed especially to be suitable for applied research, open for everyone to use and modify, well-documented, and suited to all common computer environments. During the years, GECCO has proven to be relevant for several research groups that have adopted GECCO for research experiments.

Apart from this description and the appended paper, GECCO is described in Section 3.6 where a description of recent advances along with architecture schemas, a screenshot and a photo from a recent laboratory experiment can be found.

5.4 Paper IV: Enhanced Situation Awareness using Random Particles

Modern C2 systems present the current view of the situation through a situation picture that is being built up from fused sensor data. However, merely presenting a comprehensible description of the situation does not give the commander complete awareness of the development of a situation. This article presents a generic tool for prediction of forthcoming troop movements. The technique is similar to particle filtering, a method used for approximate inference in dynamic Bayesian networks by using a combination of tracking and prediction.

The prediction tool has been implemented and installed into an existent electronic warfare system. The tool makes use of the system's geographic information system to extract geographic properties and calculate troop velocities in the terrain which is, in turn, being used for the construction of the tool's transition model. Finally, the result is presented together with the situation picture.

The prediction tool has been evaluated in field tests performed in cooperation with the Swedish Armed Forces in an exercise in Sweden during the spring of

2005. Officers and operators of the electronic warfare system were interviewed and exposed to the tool. Reactions were positive and prediction of future troop movements was considered to be interesting for short-term tactical C2.

5.5 Paper V: A Toolbox for Multi-Attribute Decision-Making

The paper discusses simulation based decision support tools for decision-making with a particular focus on the tool's relation to the overall decision process that the commander is involved in. The theoretical underpinnings focus on several contradicting decision-making principles originating from cognitive psychology.

Pure utility theory fails to address certain decision problems. Such problems arise when there exist conditions on the attributes reflecting that they are more or less worth in different situations. One example could be that a certain attribute is not worth anything before a certain utility value threshold is reached, e.g., an apartment is deemed worthless if it is not located within 10 kilometers from work. Another example could be that the utility of a certain attribute diminishes when its value improves, i.e., it is more useful to come a kilometer closer to work for a person that lives two kilometers away than for a person that lives one kilometer away. Many different principles that to some extent account for these shortcomings are reviewed in the paper. For example, the *conjunctive rule* states that the utility in each attribute should not be allowed to be below a certain threshold value, while the *lexicographic rule* states that the strategy that is best in the most important attribute should be chosen. Obviously, these criteria cannot be combined in general so that all of them are fulfilled at the same time. The paper therefore uses the criteria as a basis for a preference function to be used for multi-attribute decision-making. The preference function combines the given rules so that they are all satisfied to some extent.

Finally, by utilizing the proposed preference function a tool supporting the described decision-making process is envisioned.

The contributions of the paper are 1) a description of an agent-based decision process, 2) a non-linear multi-attribute utility function intended to better fit prevailing cognitive decision-making models than traditional linear utility functions do, 3) a toolbox concept that, based on the proposed utility function, uses embedded simulation and evolutionary learning to evolve strategies and to support decision-making.

5.6 Paper VI: An Information Fusion Game Component

Prediction of future course of events is a necessary ingredient in tomorrow's C2 centers, which is also envisioned in higher levels of, e.g., the JDL data fusion model, see Section 2.5, where awareness of the development of a situation is deemed crucial for providing a complete and comprehensible situation picture. To cope with gaming situations, i.e., situations where commanders' make decisions based on other commanders' reasoning about one's own reasoning, traditional AI methods for inference need to be extended with algorithms stemming from game theory.

In this paper, the two views presented in Chapter 4, i.e., probabilistic expert systems and game theory, are brought together in one homogenous architecture coined the "information fusion game component". The paper outlines a schematic model using influence diagrams to obtain parameters for a description of the situation in the form of a Bayesian game. The result from the game is a description of equilibrium strategies for participants that can be incorporated in the influence diagram to form a Bayesian network description of the situation and its development, hence, changing decision nodes to chance nodes. The concept of a Bayesian game makes it possible for a commander to incorporate prior beliefs regarding his opponents and seems to be a good choice for representing realistic situations. Moreover, the underlying influence diagrams make it possible to derive the game utilities using an information structure that provide means for representation of realistic, potentially large and complex, situation descriptions.

An interesting point in the paper is the assumption of consistent beliefs that is adopted throughout the work. It encompasses a somewhat philosophical question regarding the viewpoint taken for game-theoretic analyses, namely stating that a Bayesian game with consistent beliefs is a game where the player's belief, conditional on his type, about other players' types are all derivable from a global distribution over all players' types by conditioning, i.e., $p_i(t_{-i} \mid t_i) = p(t_{-i} \mid t_i)$. The assumption of consistent beliefs is both required and natural for most applications; it simply means one should model the opponents using all currently available information. Although game theory states that we should solve a game for all players at the same time, the solution is still being obtained from one particular decision-maker's view of the situation. Therefore, consistent versus inconsistent beliefs becomes more of a philosophical, but still interesting, question. Hence, the game component assumes consistent beliefs.

5.7 Paper VII: Refinements of the Command and Control Game Component

Taking the envisioned game component and a discussion regarding the game size that can be anticipated in C2 situations as a starting point, this paper discusses the computational bottlenecks and ambiguities that exist when computing optimal game-theoretic solutions in the form of Nash equilibria.

Envisioning Nash equilibria as the optimal solution concept to strategic reasoning problems poses concerns in several respects. Game-theoretic methods for prediction are, in most cases, intractable for the generic case. The best known method, the Lemke-Howson algorithm, yields exponential worst case running time and does not find all solutions. Moreover, the fragile equilibria concept makes algorithms hard to implement in practice due to numerical instability. Lastly, there sometimes exist multiple solutions that are not possible to rank versus each other in accordance to some rationally sound scheme. However, the Nash equilibria remain the rationally sound solution and, moreover, serious alternatives are non-existent.

To get an understanding of the games and their properties that the C2 game component gives rise to, computer experiments using a state-of-the-art equilibrium computation package are presented. The computer simulations, based on games sampled according to the game component idea, indicate that computation of optimal solutions seems to be tractable in reasonably sized C2 decision problems. Moreover, despite the intractability of finding all optimal solutions there exist reasonably fast algorithms that often finds all, or nearly all, solutions.

Chapter 6

Concluding Remarks

We have studied information and uncertainty management in command and control (C2) with a particular focus on handling strategic interaction in C2 systems. The area has been approached from several directions resulting in both applied and theoretical results. Applied research directions include software development, prototype implementation, and end user testing. Theoretical research directions include algorithm development, specification of information handling infrastructure, management of multiple attributes, and improvement of decision-theoretic mechanisms to account for realistic settings. We review the results below by discussing the thesis' overall perspective in Section 6.1, the applied results in Section 6.2, the theoretical results in Section 6.3, and suggestions for future undertakings in Section 6.4.

6.1 The Command and Control Gaming Perspective

In this thesis we have taken a broad perspective on gaming in order to incorporate different aspects of strategic interaction. Without reducing the value of the papers supporting the thesis, we believe that the *overall* contribution of the thesis is this main theme: a gaming perspective on command and control.

It should, however, be pointed out that although the game is still there, the gaming perspective is not always appropriate. To be more precise, the strategic situations of interest are situations in which an actor must anticipate what others will do and what others will infer from the actor's own actions. Adhering to our broad gaming perspective, the important features in such situations consist of *strategies* and *players*. Players choose strategies depending on the information they have, and how they rate the desirability, or *utility*, of resulting outcomes. In game theory, these aspects are formalized together with precise rules for the order in which players choose strategies. It should be noted that game-theoretic situations are not (at all) different from other situations. That is, game theory is not applicable solely to a certain kind of situation, but merely has a clear focus on the problem of how rational players would play given a formalized model of the situation.

As indicated, we do not distinguish significantly between "strategic interaction," "gaming," or "simulation," which we regard as the act of executing an underlying model. Nor do we make a clear distinction between "situations," "games," or "models," which we regard as the underlying model that can, in one way or another, be executed. It is part of our governing ideas and conviction to mix these things up a bit. Admittedly, however, although the terms may be used interchangeably they are still different in that they target different kinds of end usage. It follows that the words can be expressed in terms of each other, e.g., a simulation is a game that can be easily manipulated to test "what if?" questions whilst a game is a playable simulation. Hence, the absolute line between decision theory, information fusion, simulation, gaming, and experimentation cannot be drawn; instead, the construction of C2 systems should be seen as a joint effort with several contributors. In particular, the view of game theory as an "advanced mathematical toolbox" is detrimental and poses limitations on a wider usage of the theory.

6.2 Tools in Support of Gaming and Simulation

The U.S. military maintains an inventory of nearly 600 different wargames, simulations and models (Dunnigan, 2000, p. 320). These are used for simulation, training, and for wargaming. That is, games and simulations are primarily used for two purposes: to obtain information or to gain experience. A simulation means asking a "what if?" question, i.e., to set some parameters, run the simulation, and see what happens. The piece of obtained information can be used for evaluation of new technology or as input to, e.g., functions within a C2 system. Within the scope of this thesis we have outlined and implemented an example tool for analyzing decisions using the "what if?" methodology in Paper II. Training, on the other hand, is an example of usage directed towards gaining experience. Wargaming is somewhere in between and can be used both for experience and for information gathering. That is, the military commander involved in playing the game may benefit from the gained experience while a researcher may benefit from the obtained empirical data that can be used for, e.g., assessment of new technology or the enhancement of staff procedures. In Paper III in this thesis, we have presented GECCO, which we believe targets primarily this latter form of gaming characteristic. We consider GECCO being a novel research tool for microworld gaming where the novelty lies in that it is provided as open source software and in that it is genericly suitable for a wide variety of spatial microworlds.

The random particles presented in Paper IV constitute a more applied contribution which, based on a standard inference algorithm, helps the commander to visualize possible opponent troop movements. Reactions from a field evaluation resulted in positive reactions regarding a tool for prediction of future troop movements. Integration requires more research regarding how such a tool should be integrated and regarding the impact such a tool has on the human operator's actual decision-making.

The view that practical gaming and game theory are correlated is influenced from the area of C2 decision-making where this correlation is needed due to its inherent need to take conflict and, at the same time, realism into account. To create C2 systems that take strategic interaction into account, these two worlds must be bridged.

6.3 Command and Control Game-Theoretic Modeling

Gaming is an inherent part of C2 decision-making. For commanders wishing to optimize their decisions in complex multi-agent environments, understanding the rules of the game is often the same thing as understanding the decision problem itself. This is what long-term planning is all about: to take into account one's knowledge or expectation of other decision-makers' behavior to form a systematic description of the outcomes that may emerge.

Developers of tomorrow's C2 centers are facing numerous problems related to the large amount of available information obtained from various sources. On a lower level, uncertain reports from different sensors need to be fused into comprehensible information. On a higher level, representation and management of the aggregated information will be the main task, with prediction of future course of events being the ultimate goal. As indicated, information handling will be an essential issue. Consequently, Paper I proposes precision, quality, and utility to be three important measures of information to be used for information handling in C2 and suggests how these concepts can be used and interpreted. As a result, we state that the concept of situation awareness should comprise awareness of information in terms of its usefulness, to what extent it can be trusted, and the benefit of gaining more information.

Because of its limited capability to deal with real-world friction, decision theory and game theory in its purest classical form cannot be used on its own for handling C2 decision-making. However, at some point decision-making models need to be made concrete and explicit to be incorporated in C2 systems. Following this line of reasoning, Paper V proposes a multi-attribute utility function intended to account for a number of contradicting decision-making principles originating from cognitive psychology.

Traditional agent modeling techniques do not capture situations where commanders make decisions based on other commanders' reasoning about one's own reasoning. To cope with this problem, Paper VI proposes a decision support tool for C2 situation awareness enhancements based on game theory for inference and coupled with traditional AI methods for uncertainty modeling. By extending readily available and accepted single-agent reasoning engines in the form of Bayesian networks and its extensions into a "game component" we have constructed an architecture that is envisioned to be well suited for C2 reasoning. Our ideas assume that each decision-maker is rational in the sense that he is aware of his alternatives, forms expectations about any unknowns, has clear preferences, and chooses his action deliberately after some process of optimization. The assumption of rationality is not undisputed, being under perpetual attack by experimental psychologists who point out limits to its application (Osborne and Rubinstein, 1994, p. 5). However, the use of Bayesian games can to a large extent compensate for irrational behavior by letting the commander maintain a belief over several, possibly "irrational," opponent models. Hence, we claim that irrationality should be modeled in a rational manner and, likewise, that the solution should be interpreted in light of this rationally modeled irrationality. That is, we cannot reach out and "grab the irrational," but we can gain a higher state of awareness and make more well-informed decisions by incorporating the irrational using a rational model. We believe our envisioned game component provides a means for commanders wishing to make their mental models concrete and explicit which seems like a reasonable requirement in C2 decision-making. That is, commanders' mental models must be confronted and inspected to make it possible to create more elaborated models and to be able to grasp an increasingly complex situation.

Paper VII examines the characteristics of the game component with respect to computational tractability in laboratory settings, but the conditions under which the game component will be truly effective and how it should be integrated in actual C2 systems remain to be analyzed. Such studies will, in turn, provide basic data for further research regarding computational tractability, equilibrium appropriateness, and so forth.

6.4 Directions for Future Work

More realistic and complex scenarios are required to obtain adequate understanding of the difficulties and the possibilities that a game-theoretic solution yields. The solution concept in the form of mixed strategy Nash equilibria rests on well-established assumptions, but unfortunately the solution itself is often non-trivial and therefore some level of understanding of the underlying concepts is required from the commander. Further research is needed regarding how to establish such understanding. Our belief is that the development of game-theoretic tools must be made in parallel with the development of planning methods for C2 decision-making which will facilitate in establishing understanding as well as ensuring that the result matches the actual decision-making process. We believe the latter to be an important usability aspect that needs to be considered in further research and development.

Regarding the game component, two research directions are on our current research agenda. First, the relation between the uppermost chance node, i.e., the move of nature, and the area of robust Bayesian analysis should be investigated. Our hope is that one can specify a "robust game component". Second, the multi-agent influence diagram (MAID) game representation presented by Koller and Milch

(2003) provides, as it seems, the possibility to write down an architecture taking consecutive or hierarchical decisions into account. The MAID representation makes it possible to decompose certain games into a set of interacting smaller games, which can be solved in sequence. The decomposition leads to substantial savings in the computational cost of finding Nash equilibria for games that can be satisfactorily decomposed. We believe that certain classes of games resulting from realistic settings and involving consecutive or hierarchical decisions can be handled efficiently using the MAID representation.

The GECCO architecture along with its siblings should be developed further with emphasis on the architecture's generic properties and its availability. A prosperous community sharing scenarios and experiences would indefinitely boost GECCO's performance and availability.

Bibliography

Simon Ahlberg, Pontus Hörling, Karsten Jöred, Christian Mårtenson, Göran Neider, Johan Schubert, Hedvig Sidenbladh, Pontus Svenson, Per Svensson, Katarina Undén, and Johan Walter. 2007. An information fusion demonstrator for tactical intelligence processing in network-based defense. *Information Fusion*, 8(1). In press.

David S. Alberts, John J. Garstka, and Frederick P. Stein. 1999. *Network Centric Warfare: Developing and Leveraging Information Superiority*, second edition. CCRP, DoD, Washington, District of Columbia.

David S. Alberts and Richard E. Hayes, editors. 2002. *Code of Best Practice for Experimentation.* Information Age Transformation Series. CCRP, DoD, Washington, District of Columbia.

David S. Alberts and Richard E. Hayes. 2006. *Understanding Command and Control.* The Future of Command and Control. CCRP, DoD, Washington, District of Columbia.

Stefan Arnborg, Henrik Artman, Joel Brynielsson, and Klas Wallenius. 2000. Information awareness in command and control: Precision, quality, utility. In *Proceedings of the Third International Conference on Information Fusion (FUSION 2000)*, pages ThB1/25–32, Paris, France, July 2000. Paper I in this thesis.

Henrik Artman. 1999. *Fördelade kunskapsprocesser i ledningscentraler vid nödsituationer – koordination och situationsmedvetenhet* [Distributed Cognition in Control Rooms for Emergency Situations – Coordination and Situation Awareness]. PhD thesis, Department of Communication Studies, Linköping University, Sweden.

Thomas Bayes. 1763. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53:370–418.

Mark Bedworth and Jane O'Brien. 2000. The omnibus model: A new model of data fusion? *IEEE Aerospace and Electronic Systems Magazine*, 15(4):30–36.

Bruce Berkowitz. 2003. *The New Face of War: How War Will Be Fought in the 21st Century.* Free Press, New York.

Samuel S. Blackman and Robert Popoli. 1999. *Design and Analysis of Modern Tracking Systems*. Artech House, Norwood, Massachusetts.

Alfred A. Bove. 1998. Risk of decompression sickness with patent foramen ovale. *Undersea & Hyperbaric Medicine*, 25(3):175–178.

Berndt Brehmer. 1992. Dynamic decision making: Human control of complex systems. *Acta Psychologica*, 81(3):211–241.

Berndt Brehmer. 2000. Dynamic decision making in command and control. In Carol McCann and Ross Pigeau, editors, *The Human in Command: Exploring the Modern Military Experience*, chapter 16, pages 233–248. Kluwer Academic/Plenum Publishers, New York.

Berndt Brehmer. 2004. Some reflections on microworld research. In Samuel G. Schiflett, Linda R. Elliott, Eduardo Salas, and Michael D. Coovert, editors, *Scaled Worlds: Development, Validation and Applications*, chapter 2, pages 22–36. Ashgate.

Berndt Brehmer and Robert Allard. 1991. Dynamic decision making: The effects of task complexity and feedback delay. In Jens Rasmussen, Berndt Brehmer, and Jacques Leplat, editors, *Distributed Decision Making: Cognitive Models for Cooperative Work*, chapter 16, pages 319–334. John Wiley & Sons.

Berndt Brehmer and Dietrich Dörner. 1993. Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior*, 9:171–184.

Berndt Brehmer and Peter Svenmarck. 1995. Distributed decision making in dynamic environments: Time scales and architectures of decision making. In Jean-Paul Caverni, Maya Bar-Hillel, F. Hutton Barron, and Helmut Jungermann, editors, *Contributions to Decision Making – I*. Elsevier Science.

Joel Brynielsson. Using AI and games for decision support in command and control. *Decision Support Systems*. Accepted for publication.

Joel Brynielsson. 2002. A decision–theoretic framework using rational agency. In *Proceedings of the 11th Conference on Computer-Generated Forces and Behavioral Representation*, number 02–CGF–047, pages 459–463, Orlando, Florida, May 2002.

Joel Brynielsson. 2004. Game-theoretic reasoning in command and control. In *Proceedings of the 15th Mini-EURO Conference: Managing Uncertainty in Decision Support Models (MUDSM 2004)*, Coimbra, Portugal, September 2004.

Joel Brynielsson and Stefan Arnborg. An information fusion game component. *Journal of Advances in Information Fusion*. Accepted for publication. Paper VI in this thesis.

Joel Brynielsson and Stefan Arnborg. 2004. Bayesian games for threat prediction and situation analysis. In Per Svensson and Johan Schubert, editors, *Proceedings of the Seventh International Conference on Information Fusion (FUSION 2004)*, volume 2, pages 1125–1132, Stockholm, Sweden, June 28–July 1, 2004.

Joel Brynielsson and Stefan Arnborg. 2005. Refinements of the command and control game component. In *Proceedings of the Eighth International Conference on Information Fusion (FUSION 2005)*, Philadelphia, Pennsylvania, July 2005. Paper VII in this thesis.

Joel Brynielsson, Mattias Engblom, Robert Franzén, Jonas Nordh, and Lennart Voigt. 2005. Enhanced situation awareness using random particles. In *Proceedings of the Tenth International Command and Control Research and Technology Symposium (ICCRTS)*, McLean, Virginia, June 2005. Paper IV in this thesis.

Joel Brynielsson and Rego Granlund. 2001. Assistance in decision making: Decision help and decision analysis. In *Proceedings of the Sixth International Command and Control Research and Technology Symposium (ICCRTS)*, Annapolis, Maryland, June 2001. Paper II in this thesis.

Joel Brynielsson and Klas Wallenius. 2001. Game environment for command and control operations (GECCO). In *Proceedings of the First International Workshop on Cognitive Research With Microworlds*, pages 85–95, Granada, Spain, November 2001. Paper III in this thesis.

Joel Brynielsson and Klas Wallenius. 2003. A toolbox for multi-attribute decision-making. Technical Report TRITA–NA–0307, Department of Numerical Analysis and Computer Science, Royal Institute of Technology, Stockholm, Sweden. Paper V in this thesis.

Colin F. Camerer. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction.* The Roundtable Series in Behavioral Economics. Princeton University Press, New Jersey.

Vašek Chvátal. 1983. *Linear Programming.* W. H. Freeman and Company, New York.

Carl von Clausewitz. 1976. *On War* (translated and edited by Michael Howard and Peter Paret), 1984 edition. Princeton University Press, New Jersey. Originally published 1832.

Thomas P. Coakley. 1991. *Command and Control for War and Peace.* National Defense University Press, Washington, District of Columbia.

Gregory F. Cooper. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2–3):393–405.

Richard W. Cottle, Jong-Shi Pang, and Richard E. Stone. 1992. *The Linear Complementarity Problem.* Academic Press.

Robert G. Cowell, A. Philip Dawid, Steffen L. Lauritzen, and David J. Spiegelhalter. 1999. *Probabilistic Networks and Expert Systems.* Statistics for Engineering and Information Science. Springer-Verlag, New York.

Paul Dagum and Michael Luby. 1993. Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artificial Intelligence*, 60(1):141–153.

James F. Dunnigan. 2000. *Wargames Handbook: How to Play and Design Commercial and Professional Wargames*, third edition. Writers Club Press, Lincoln, Nebraska.

Erik Durschmied. 1999. *The Hinge Factor: How Chance and Stupidity Have Changed History.* Coronet Books, London, United Kingdom.

Dietrich Dörner, Heinz W. Kreuzig, Franz Reither, and Thea Stäudel, editors. 1983. *Lohhausen. Vom Umgang mit Unbestimmtheit und Komplexität* [Lohhausen. On dealing with uncertainty and complexity]. Hans Huber, Bern, Switzerland.

Dietrich Dörner, Thea Stäudel, and Stefan Strohschneider. 1986. *MORO: Programmdokumentation.* Memorandum No. 23. Lehrstuhl für Psychologie II, University of Bamberg, Germany.

Ward Edwards. 1962. Dynamic decision theory and probabilistic information processing. *Human Factors*, 4:59–73.

Fredrik Elg. 2002. *Ett dynamiskt perspektiv på individuella skillnader av heuristisk kompetens, intelligens, mentala modeller, mål och konfidens i kontroll av mikrovärlden Moro* [A Dynamic Perspective on Individual Differences in Heuristic Competence, Intelligence, Mental Models, Goals and Confidence in Control of the Microworld Moro]. Licentiate thesis, Department of Computer and Information Science, Linköping University, Sweden.

Mica R. Endsley. 1988. Design and evaluation for situation awareness enhancement. In *Proceedings of the Human Factors Society 32nd Annual Meeting*, pages 97–101, Santa Monica, California.

Mica R. Endsley. 1995. Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64.

Mica R. Endsley. 2000. Theoretical underpinnings of situation awareness: A critical review. In Mica R. Endsley and Daniel J. Garland, editors, *Situation Awareness Analysis and Measurement*, pages 3–32. Lawrence Erlbaum Associates, Inc., Mahwah, New Jersey.

Anders Frank and Robert Virding. 2003. The “new” world of simulation – digital games and gaming. In *Proceedings of the 2003 European Simulation Interoperability Workshop*, number 03E–SIW–073, pages 492–497, Stockholm, Sweden, June 2003.

Itzhak Gilboa and Eitan Zemel. 1989. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93.

Piotr J. Gmytrasiewicz and Edmund H. Durfee. 2000. Rational coordination in multi-agent environments. *Autonomous Agents and Multi-Agent Systems*, 3(4): 319–350.

Rego Granlund. 1997. *C^3Fire – A Microworld Supporting Emergency Management Training*. Licentiate thesis, Department of Computer and Information Science, Linköping University, Sweden.

Abe Greenberg. 1981. An outline of wargaming. *Naval War College Review*.

John C. Harsanyi. 1967–1968. Games with incomplete information played by “Bayesian” players. *Management Science*, 14(3,5,7):159–182, 320–334, 486–502.

Ronald A. Howard and James E. Matheson. 1984. Influence diagrams. In Ronald A. Howard and James E. Matheson, editors, *Readings on the Principles and Applications of Decision Analysis*, pages 721–762. Strategic Decisions Group, Menlo Park, California.

Qi Huang, Jenny Hållmats, Klas Wallenius, and Joel Brynielsson. 2003. Simulation-based decision support for command and control in joint operations. In *Proceedings of the 2003 European Simulation Interoperability Workshop*, number 03E–SIW–091, pages 591–599, Stockholm, Sweden, June 2003.

Edwin T. Jaynes. 2003. *Probability Theory: The Logic of Science*. Cambridge University Press, United Kingdom.

Eva Jensen and Berndt Brehmer. 2003. Understanding and control of a simple dynamic system. *System Dynamics Review*, 19(2):119–137.

Finn V. Jensen. 2001. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York.

Björn Johansson. 2005. *Joint control in dynamic situations*. PhD thesis, Department of Computer and Information Science, Linköping University, Sweden.

Ronnie Johansson and Robert Suzić. 2005. Particle filter-based information acquisition for robust plan recognition. In *Proceedings of the Eighth International Conference on Information Fusion (FUSION 2005)*, Philadelphia, Pennsylvania, July 2005.

Stuart E. Johnson and Martin C. Libicki, editors. 1995. *Dominant Battlespace Knowledge*. National Defense University Press, Washington, District of Columbia.

Michael I. Jordan, editor. 1999. *Learning in Graphical Models*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, Massachusetts.

Paul R. Kleindorfer, Howard C. Kunreuther, and Paul J. H. Schoemaker. 1993. *Decision Sciences: An Integrative Perspective*. Cambridge University Press, United Kingdom.

Janeen Klinger. 2006. The social science of Carl von Clausewitz. *Parameters: US Army War College Quarterly*, XXXVI(1):79–89.

Daphne Koller and Brian Milch. 2003. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45(1):181–221.

Daphne Koller and Avi Pfeffer. 1997. Representations and solutions for game-theoretic problems. *Artificial Intelligence*, 94(1):167–215.

Andrey N. Kolmogorov. 1933. *Grundbegriffe der Wahrscheinlichkeitrechnung* [Foundations of the Theory of Probability]. Springer-Verlag, Berlin, Germany.

Harold W. Kuhn. 1953. Extensive games and the problem of information. In Harold W. Kuhn and Albert W. Tucker, editors, *Contributions to the Theory of Games II*, pages 193–216. Princeton University Press, New Jersey.

Jan Kuylenstierna, Joacim Rydmark, and Hans Sandström. 2004. Some research results obtained with DKE: A dynamic war-game for experiments. In *Proceedings of the Ninth International Command and Control Research and Technology Symposium (ICCRTS)*, Copenhagen, Denmark, September 2004.

Averill M. Law and W. David Kelton. 1991. *Simulation Modeling and Analysis*, second edition. McGraw-Hill Series in Industrial Engineering and Management Science.

Carlton E. Lemke and Joseph T. Howson, Jr. 1964. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2): 413–423.

James Llinas, Christopher Bowman, Galina Rogova, Alan Steinberg, Ed Waltz, and Franklin E. White, Jr. 2004. Revisiting the JDL data fusion model II. In Per Svensson and Johan Schubert, editors, *Proceedings of the Seventh International Conference on Information Fusion (FUSION 2004)*, volume 2, pages 1218–1230, Stockholm, Sweden, June 28–July 1, 2004.

R. Duncan Luce and Howard Raiffa. 1957. *Games and Decisions*. John Wiley & Sons, New York.

Leif Løvborg and Berndt Brehmer. 1991. NEWFIRE – A flexible system for running simulated fire fighting experiments. Report Risø–M–2953, Risø National Laboratory, Denmark.

Roger B. Myerson. 1991. *Game Theory: Analysis of Conflict.* Harvard University Press.

John F. Nash. 1951. Non-cooperative games. *Annals of Mathematics*, 2(54):286–295.

John von Neumann. 1928. Zur Theorie der Gesellschaftsspiele [On the theory of games of strategy]. *Mathematische Annalen*, 100:295–320.

John von Neumann and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior.* Princeton University Press, New Jersey.

Mary Omodei and Alex Wearing. 1995. The fire chief microworld generating program: An illustration of computer-simulated microworlds as an experimental paradigm for studying complex decision making behaviour. *Behavior Research Methods, Instruments, & Computers*, 27:303–316.

Martin J. Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory.* MIT Press, Cambridge, Massachusetts.

Judea Pearl. 1982. Reverend Bayes on inference engines: A distributed hierarchical approach. In David L. Waltz, editor, *Proceedings of the National Conference on Artificial Intelligence (AAAI-82)*, pages 133–136, Pittsburgh, Pennsylvania, August 1982.

Judea Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann Publishers, Inc., San Mateo, California.

Peter P. Perla. 1990. *The Art of Wargaming: A Guide for Professionals and Hobbyists.* Naval Institute Press, Annapolis, Maryland.

Peter P. Perla and Julia J. Loughran. 2003. Using gaming and agent technology to explore C2. In *Proceedings of the Eighth International Command and Control Research and Technology Symposium (ICCRTS)*, Washington, District of Columbia, June 2003.

Avi Pfeffer. 1999. *Probabilistic Reasoning for Complex Systems.* PhD thesis, Stanford University, California.

Howard Raiffa. 1968. *Decision Analysis: Introductory Lectures on Choices under Uncertainty.* Addison–Wesley, Reading, Massachusetts.

Jens Rasmussen, Berndt Brehmer, and Jacques Leplat, editors. 1991. *Distributed Decision Making: Cognitive Models for Cooperative Work.* John Wiley & Sons.

Georgios Rigas. 2000. *On the relationship between psychometric intelligence and decision making in dynamic systems.* PhD thesis, Uppsala University, Department of Psychology, Uppsala University, Sweden.

Stuart J. Russell and Peter Norvig. 2003. *Artifical Intelligence: A Modern Approach*, second edition. Prentice Hall, Upper Saddle River, New Jersey.

Joacim Rydmark. 2002. *Beslut under osäkerhet: En experimentell mikrovärldsstudie av metoder för att hantera osäkerhet vid ledning och beslutsfattande i komplexa och dynamiska miljöer* [Decision-Making Under Uncertainty: An Experimental Microworld Study of Methods for Handling Uncertainty in Command and Control and Decision-Making in Complex and Dynamic Environments]. Licentiate thesis, School of Public Administration, Göteborg University, Sweden.

John Salerno. 2002. Information fusion: A high-level architecture overview. In *Proceedings of the Fifth International Conference on Information Fusion (FUSION 2002)*, Annapolis, Maryland, June 2002.

Glenn Shafer. 1996. *Probabilistic Expert Systems.* Number 67 in CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania.

Edward Allen Smith, Jr. 2002. *Effects Based Operations: Applying Network Centric Warfare in Peace, Crisis, and War.* Information Age Transformation Series. CCRP, DoD, Washington, District of Columbia.

Alan Steinberg, Christopher Bowman, and Franklin E. White, Jr. 1999. Revisions to the JDL data fusion model. In *Proceedings of SPIE AeroSense (Sensor Fusion: Architectures, Algorithms, and Applications III)*, volume 3719, pages 430–441, Orlando, Florida, April 1999.

Bernhard von Stengel. 2002. Computing equilibria for two-person games. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 3 of *Handbooks in Economics*, chapter 45, pages 1723–1759. Elsevier Science.

Sun Tzu. 1994. *The Art of War* (translated by Ralph D. Sawyer). Barnes & Noble.

Edward J. Villacres and Christopher Bassford. 1995. Reclaiming the Clausewitzian trinity. *Parameters: US Army War College Quarterly*, XXV(3):9–19.

Klas Wallenius. 2002. A generic model of management and command and control. In *Proceedings of the Seventh International Command and Control Research and Technology Symposium (ICCRTS)*, Quebec City, Canada, September 2002.

Klas Wallenius. 2004. DISCCO: Generic support for C2 and decision-making. In *2004 Command and Control Research and Technology Symposium (CCRTS)*, San Diego, California, June 2004.

Klas Wallenius. 2005. *Generic Support for Decision-Making in Effects-Based Management of Operations*. PhD thesis, School of Computer Science and Communication, Royal Institute of Technology, Stockholm, Sweden.

A. Rodney Wellens and Deniz Ergener. 1988. The C.I.T.I.E.S. game: A computer-based situation assessment task for studying distributed decision making. *Simulation & Games*, 19(3):304–327.

Franklin E. White, Jr. 1987. Data fusion lexicon. Technical report, Joint Directors of Laboratories, Technical Panel for C3, Data Fusion Sub Panel, Naval Ocean Systems Center, San Diego, California.

Franklin E. White, Jr. 1988. A model for data fusion. In *Proceedings of the First National Symposium on Sensor Fusion*, volume 2, Orlando, Florida, April 1988.