

# Planning Under Risk and Uncertainty: Optimizing Spatial Forest Management Strategies

Nicklas Forsell

*Faculty of Forest Science*

*Department of Forest Resource Management*

*Umeå*

Doctoral Thesis

Swedish University of Agricultural Sciences

Umeå 2009

Acta Universitatis agriculturae Sueciae

2009:39

ISSN 1652-6880

ISBN 978-91-86195-86-1

© 2009 Nicklas Forsell, Umeå

Print: Arkitektkopia 2009

# Planning under risk and uncertainty: Optimizing spatial forest management strategies

## Abstract

This thesis concentrates on the optimization of large-scale management policies under conditions of risk and uncertainty.

In paper I, we address the problem of solving large-scale spatial and temporal natural resource management problems. To model these types of problems, the framework of graph-based Markov decision processes (GMDPs) can be used. Two algorithms for computation of high-quality management policies are presented: the first is based on approximate linear programming (ALP) and the second is based on mean-field approximation and approximate policy iteration (MF-API). The applicability and efficiency of the algorithms were demonstrated by their ability to compute near-optimal management policies for two large-scale management problems. It was concluded that the two algorithms compute policies of similar quality. However, the MF-API algorithm should be used when both the policy and the expected value of the computed policy are required, while the ALP algorithm may be preferred when only the policy is required.

In paper II, a number of reinforcement learning algorithms are presented that can be used to compute management policies for GMDPs when the transition function can only be simulated because its explicit formulation is unknown. Studies of the efficiency of the algorithms for three management problems led us to conclude that some of these algorithms were able to compute near-optimal management policies.

In paper III, we used the GMDP framework to optimize long-term forestry management policies under stochastic wind-damage events. The model was demonstrated by a case study of an estate consisting of 1,200 ha of forest land, divided into 623 stands. We concluded that managing the estate according to the risk of wind damage increased the expected net present value (NPV) of the whole estate only slightly, less than 2%, under different wind-risk assumptions. Most of the stands were managed in the same manner as when the risk of wind damage was not considered. However, the analysis rests on properties of the model that need to be refined before definite conclusions can be drawn.

*Keywords:* Forest management, planning under risk and uncertainty, spatial processes, factored markov decision processes, collaborative multiagent markov decision processes, graphical models, optimization, reinforcement learning.

*Author's address:* Nicklas Forsell, Department of Forest Resource Management, SLU Skogsmarksgränd, SE-901 83 Umeå, Sweden.

*E-mail:* Nicklas.Forsell@srh.slu.se

# Dedication

To Emilie-Anne Guerch.

# Contents

<b>List of Publications</b>	<b>7</b>
<b>1 Introduction</b>	<b>9</b>
1.1 Forest Planning	9
1.2 Forest Planning under Risk and Uncertainty	10
1.2.1 Wind damage	11
1.2.2 Fire management	13
1.2.3 Losses from insect and fungal disease	15
1.2.4 Other forest-level related problems	16
1.2.5 State of the art	18
1.3 Markov Decision Processes	18
1.3.1 Factored MDP	21
1.3.2 Collaborative multiagent MDPs	24
<b>2 Objectives and main contribution of the thesis</b>	<b>27</b>
<b>3 Summary of Papers</b>	<b>29</b>
3.1 A framework and algorithms for the local control of spatial processes (Paper I)	29
3.2 Q-learning for graph-based Markov decision processes (Paper II)	32
3.3 Management of the risk of wind damage in forestry: a graph-based Markov decision process approach (Paper III)	35
<b>4 Discussion and Conclusion</b>	<b>40</b>
4.1 Modeling	40
4.2 Efficient model-based algorithms	41
4.3 Efficient model-free algorithms	43
4.4 Multiagent coordination	43
4.5 Collaborative multiagent MDPs	44
4.6 Forest planning under risk of wind damage	44
<b>5 Future Research</b>	<b>46</b>
<b>References</b>	<b>49</b>
<b>Acknowledgements</b>	<b>58</b>



## List of Publications

This thesis is based on the work contained in the following papers, referred to by Roman numerals in the text:

- I Forsell, N., Peyrard, N., Sabbadin, R. A framework and algorithms for the local control of spatial processes (manuscript).
- II Forsell, N., Garcia, F., Sabbadin, R. Q-learning for graph-based Markov decision processes (manuscript).
- III Forsell, N., Wikström, P., Garcia, F., Sabbadin, R., Blennow, K., Eriksson, L.O. Management of the risk of wind damage in forestry: a graph-based Markov decision process approach. *Annals of Operations Research*, Published online 05 February 2009.

Papers III is reproduced with the permission of the publisher.

## Abbreviations

LP	Linear programming
ALP	Approximate linear programming
MF	Mean-field approximation
API	Approximate policy iteration
RL	Reinforcement learning
MDPs	Markov decision processes
GMDPs	Graph-based Markov decision processes

# 1 Introduction

## 1.1 Forest Planning

Forest managers commonly need to decide when and what types of management activities should be performed in their forest. Activities such as stand establishment, thinning, fertilization, selective cutting, and clear-cutting will ultimately determine the outcome and the trajectory the forest will follow over time. A forest manager may need a policy that optimizes the management of the forest. The management policy is optimized according to some objective decided by the decision-maker; it can include economic values, biodiversity, scenography, hunting, and so forth.

The planning problem has traditionally been divided into strategic, tactical, and operational planning (Church, 2007; Epstein *et al.*, 2007; Davis *et al.*, 2001; Church *et al.*, 2000; Church *et al.*, 1994). These problems differ essentially in the length of their planning horizon. Strategic planning considers broad-scale planning over a long period of time and over a large landscape. Typically, time is aggregated into five- or ten-year time periods, and a planning horizon ranging from 50, 100, or 150 years or even more is commonly used. Strategic forest planning was traditionally designed to maximize sustained and preferably constant harvest volume flow, but has now been extended to focus also on other issues such as nature conservation, biodiversity, wildlife preservation, aesthetics, recreation, and habitat structure of the forest. The aim of tactical planning is to identify how the specifications set during the first time periods of the strategic forest policy can be met. The time periods used by the strategic planning are split up into shorter time periods of usually one year, and questions concerning roads, budget, logging, development of facilities, and the amount of outside timber to be purchased are considered. A planning horizon of 5 to 10 years is

generally used. Tactical planning thereby helps to form a bridge between the strategic policy and the operational policy. Operational planning deals with the actual operations that need to be carried out in silviculture and harvesting. Typically, weekly or monthly plans concerning cutting units, location and use of harvesting machinery, logging transportation, and tree stem bucking patterns are considered. In this thesis, we will concentrate on strategic planning, tactical and operational planning will not be discussed further.

Another way of classifying planning problems is to divide them into stand-level and forest-level planning problems. A stand-level model deals with the decisions concerning a single stand whereas a forest-level model deals with the decisions concerning a number of stands at the same time. In some forest-level models, the stands are aggregated into a few strata, while in other models explicit representation of each stand is kept.

## 1.2 Forest Planning under Risk and Uncertainty

As strategic planning considers a long planning period, considerable uncertainties concerning the future state of the forest and the effect of different management activities have to be dealt with. Some of the uncertainties are due to natural disturbances such as forest fire, wind damage, disease, and insect attacks. These natural disturbances can be viewed as stochastic, or random, events as we cannot precisely determine if, when, or where they will occur. Other sources of uncertainty are changes in market prices, interest rates, technological developments, harvesting technologies, and social and political change. Uncertainty concerning our understanding of the forest ecosystem, the development of forest ecosystems, and incomplete data on the current state of the forest, also has to be considered.

Traditionally, forest managers have dealt with uncertainties by adjusting yield tables to reflect expected or average values estimated over long time periods. An average amount of damage due to, for example, wind is computed or estimated and the yield tables are reduced by this amount. However, as stochastic and deterministic processes have different effects on the development of the forest ecosystem (McCarthy & Burgman, 1995), and as the natural disturbances and other sources of uncertainties can have a large influence on the forest, questions concerning the incorporation of uncertainties have been raised. Can we improve the forest management plan by considering the uncertainties when optimization of the strategic plan? Theoretically, if the uncertainties are ignored, then the selected management policy is likely to be suboptimal. There are, however, a number of

difficulties in incorporating stochastic uncertainties into the optimization of the strategic management policy. One difficulty is the sheer size of the problem. Strategic management often considers large forest areas. Hundreds to thousands of stands may be involved, making the models very large. Especially in the case of natural disturbances, it is common that the intensity and severity vary over the forest. This limits the possibility of aggregating the forest into a few strata and thereby reducing the size of the model. The most demanding aspect, though, is that a number of uncertainties have an inherent spatial structure that must be considered. For example, is the probability of a stand being damaged by wind is dependent on the state of the stand and the spatial structure of the forest (Peltola *et al.*, 1999; Gardiner *et al.*, 1997; Peltola, 1996; Lohmander & Helles, 1987; Alexander, 1964). Another example is the spruce budworm (*Choristoneura fumiferana*), the impact of which on the forest in specific outbreaks depends on the composition of the forest at the stand and landscape level (Nealis & Régnière, 2004; MacLean *et al.*, 2001; Bergeron *et al.*, 1995).

### 1.2.1 Wind damage

Wind is a major concern in forestry in numerous parts of the world, due to the massive amount of damage it causes. In Europe alone, wind storms damage 18.5 million m<sup>3</sup> of wood per year (Schelhaas *et al.*, 2003). Storms in December 1999 felled almost 200 million m<sup>3</sup> of roundwood in western Europe (UNECE/FAO, 2000). In November 2001, over 7 million m<sup>3</sup> of timber was lost in Finland (Pellikka & Järvenpää, 2003). As it has been shown that the probability of wind damage is influenced by the silvicultural treatments (Valinger & Pettersson, 1996; Quine *et al.*, 1995; Lohmander & Helles, 1987; Persson, 1975), numerous studies have focused on the task of incorporating the probability of wind damage into the optimization of the forest management policy.

Meilby *et al.*, (2001) proposed a model that optimized management of the forest under the risk of wind damage. The model considers the spatial position of the stands and also several important aspects of wind damage such as spatial interaction between the stands and the geographical structure and orientation of the forest. The risk of wind damage is explicitly modeled, based on an empirical model (Lohmander & Helles, 1987) that takes the sheltering effect of the neighboring forest stands into account. Two versions of the model were suggested. In the first model, storm events occur independently for each stand. In the second model, a storm event during a time period affects all stands in the forest. However, as the model explicitly states all possible future scenarios for the forest, the model grows

exponentially with the number of stands, polynomially with the number of time periods in the first model, and exponentially with the number of time periods in the second. For a forest consisting of 16 stands, an approximate management policy was computed with a simulated annealing algorithm (Černý, 1985; Kirkpatrick *et al.*, 1983). The experimental result showed that as the number of stands increases, the interaction between the stands also increases. Thus, the stands should be harvested depending on the state of the stand and on the state of the neighboring stands. Furthermore, a comparison between the stand- and forest-levels showed that the local dependencies between the stands may lead to an increase in the optimal harvest age. This is due to the sheltering effect between the stands, which may reduce the probability of stands being damaged by wind. Comparisons between the first and second models showed that the two models gave similar harvesting age. As the computational complexity of the second model is higher than the first model, this is an important observation.

As the algorithm for computing the management policy for the model proposed in Meilby *et al.*, (2001) is intractable for large-scale forests, an adaptive optimization approach was suggested by Meilby *et al.*, (2003). Instead of optimizing a long-term management policy under all future trajectories of the forest, a decision concerning which stands should be harvested for the current time period is optimized. The management policy is computed with an approximate dynamic programming algorithm that, for each time period, only considers two management options per stand: should the stand be harvested during the current period or the next period? The experimental results showed that the algorithm computes management policies that perform similarly to the management policy computed by the scenario-based algorithm proposed in Meilby *et al.*, (2001). However, as the algorithm only considers the management of the stands for the current time period, a new management policy has to be computed for each time period.

Another model has been suggested by Zeng *et al.*, (2007), in which the objective of the model was to minimize the number of edges vulnerable to damage from wind while keeping a high level of timber harvest and an even timber flow over the planning horizon. The method was evaluated for typical boreal forest in Finland containing 46 ha of open terrain and 395 ha of forest divided into 266 stands. In the case study, it was shown that the number of vulnerable edges could be reduced while still satisfying the economic objectives. Contrary to the model of Meilby *et al.*, (2001), the model of Zeng *et al.*, (2007) is deterministic and does not take stochastic damage events into account. A fixed management plan is computed, and it has to be re-computed whenever damage events occur.

### 1.2.2 Fire management

Forest fires have been receiving much attention for some time due to the threat to public safety, property, natural resources, and forest ecosystems. Every year, forest fires reduce the wildland and forested areas across southern Europe, Australia, USA, and Canada. In the Mediterranean countries, 50,000 fires burn 500,000 ha of forest on average per year (Vélez, 2002). In Canada, 10,000 fires burn 2.5 million ha of wildland areas on average per year (Lee *et al.*, 2002). During an extreme fire season in 1998, almost 1,700 fires burned more than 726,000 ha of forest in the Province of Alberta, Canada. The provincial government spent 242 million on forest protection during that season (Armstrong & Cumming, 2003). As the forest is influenced by repeated forest fire events and as forest management influences the ignition and spread of forest fires, numerous studies have focused on the task of incorporating forest fires into strategic forest planning. Numerous stand-level optimization models including the risk of fire have been proposed (González *et al.*, 2006; Caulfield, 1988; Reed, 1987; Reed & Errico, 1985; Reed, 1984; Martell, 1980). The consensus of these models is that incorporation of the risk of stochastic forest fire damage reduces the optimal rotation period. The expected harvested volume will be reduced due to the loss in the immature stands, and since the stands that will be harvested are younger.

Van Wagner (1983) was one of first to propose a model for incorporation of the risk of forest fire at the forest-level. A constant proportion of the forest was burned during every time period and the effect of fire damage on the long-term equilibrium of timber supply was studied. Reed & Errico (1986) included expected losses in the forest as a result of forest fires. The forest is aggregated over a set of age classes of the trees and harvest is specified in terms of area clear-cut per age class. Each time period, a proportion of the trees in the age classes is damaged by fire. The authors presented an approach for computing the optimal annual harvest when the annual proportion of the age class that is damaged by forest fire is random for each age class. To compute an estimation of the optimal harvest policy when harvest-flow constraints are present, a deterministic formulation is used. The proportion of the age class that is damaged by forest fire is assumed to be constant over the time periods, and computed as the expected value over the time horizon. The problem is thereby deterministic, and the annual harvest amount per age class can be computed with standard linear programming (LP) methods. Their approach validates the finding on the single-stand case that the projected harvest volumes are reduced when the risk of forest fire is incorporated. Gassman (1989) formulated a smaller

version of the model of Reed & Ericco (1986) by reducing the number of time periods considered and modeling the problem as a multistage stochastic problem with a finite time horizon. For each time period, a stochastic proportion of the age classes is damaged. Boychuk & Martell (1996) further extended the version of Van Wagner (1983), into a stochastic programming model. For the first time periods, high and low proportions of the age classes are damaged by forest fire with specific probabilities. Fire loss was thus proportional over the age distribution, stochastic over the time periods, and non-spatial over the age classes. They showed that a stable timber supply should be established by the creation of a buffer stock in the forest. Furthermore, doing so increased the stability of the harvest quantity and the expected harvest quantity over the planning horizon.

All models described so far have been non-spatial. The risk of damage is modeled as independent of the spatial distribution of the forest and no spatial consideration of the occurrence of damage is considered. To incorporate spatial risk and uncertainties, most research has instead focused on evaluating different management strategies with spatial simulation models (Ryu *et al.*, 2007; González *et al.*, 2005; Peter & Nelson, 2005; Armstrong, 2004; Gustafson *et al.*, 2004; Mehta *et al.*, 2004; Mohren, 2003; Gustafson *et al.*, 2000; Shifley *et al.*, 2000; Thompson *et al.*, 2000; Johnson *et al.*, 1998). A simulation model of a forest landscape provides the planner with the means to simulate future landscape development under stochastic forest growth and stochastic natural disturbances such as fire, wind, and insects. Most of the simulation models work in general in the following manner. The input to the model is the initial state of the stands making up the forest, a goal under which the management activities should be maximized or a set of applicable management strategies to be evaluated. The goal is generally total harvest or the annual allowable cut. For each time period, the management activities are optimized according to the goal, or selected among the given management activities. The management activities are performed in the stands, after which some stochastic, and sometimes spatial, event is simulated for the forest. Note that some of the simulated models incorporate several types of risk, for example, the LANDIS model (LANdscape DIsturbance and Succession model (Scheller *et al.*, 2007; Mladenoff, 2004; Mladenoff *et al.*, 1996)), which can incorporate stochastic wind, fire, and insect damage events. After the impact of the stochastic event has been registered, the forest “grows” until the end of the time period. The advantage of these simulation models is that they are very flexible regarding how the simulation is designed. Spatial and individual stand information can be taken into account in the simulation of the next state of the stands. However, to the

best of our knowledge, there is no simulation model that optimizes the management policy under stochastic development of the forest. The simulation models commonly require the user either to specify the management policy explicitly, or it is optimized under the current state of the forest and for a deterministic development of the forest.

### 1.2.3 Losses from insect and fungal disease

Something else that has an influence on the forest and the long term yield of forests is damage due to insects and fungal disease. Annual losses of timber volume in the USA due to insects and disease have been estimated to be as high as 68 million m<sup>3</sup> (Hamel & Shade, 1987). A number of models have been put forward to incorporate the impact of insects and fungal disease into the planning process. Reed & Errico (1987) proposed a stand-level and a forest-level timber supply model that assesses the influence of infection rates on the level of wood supply. The model was proposed along the same line as in Reed & Errico (1986), and for the stand-level, the stand was subjected to the risk of infection and the risk of damage by wind. These authors showed that on the stand-level, the impact of an infection may be viewed as a reduction in the net aggregate value of a stand, and that it does not have a large effect on the optimal rotation period on the stand. At the forest-level, the problem was modeled as being deterministic in that the level of infections over the age classes was selected as the average annual rate. As the forest is aggregated over the age classes, no spatial information concerning the spread of the infections was considered. The examples given indicate that there is a large difference between the effects in the short term (over 5 to 10 years) and in the long term (200 years). The impact on the long-term timber supply was probably not great, while the short-term and local effects were found to be dramatic, due to the intensity of the outbreaks of infection.

Hof *et al.*, (1997) proposed a number of spatial and deterministic optimization approaches for forest management when a pest outbreak has occurred. The pest is modeled to spread between the stands with a fixed probability, and the optimal management is computed with an integer programming approach. Moll & Chinneck (1992) proposed a deterministic linear programming model that optimized the management policy under the risk of damage from insects and fire. The management policy was optimized for the deterministic problem in which the average value of insect and fire damage was used. The deterministic management policy was compared to stochastic management strategies, determined by iteratively computing the optimal deterministic management policy for the current state of the forest

and thereafter simulating the next state of the forest. The harvest level of the deterministic management policy was found to be close to the harvest level of the stochastic management policy; however, as the stochastic management strategies were optimized under the deterministic development of the forest, they only gave approximate solutions to the stochastic problem.

A number of simulation-based models for evaluating management strategies have also been suggested. MacLean *et al.*, (2001; 2000) created the Spruce Budworm Decision Support System (SBW DSS) that can be used to calculate the marginal wood supply benefits of different management strategies. It uses growth loss and mortality to calculate the impact on future harvest levels of different management activities such as protection and salvage. Although the SBW DSS can quantify the effect of spruce budworm on different management strategies, it cannot optimize management interventions and specify how the stand should be treated.

Hennigar *et al.*, (2007) proposed a simulated model for analysis of the short-term and long-term abilities of management protection strategies to reduce volume losses due to spruce budworm. The model was used to show that a buffer should be created in the forest for difficult years to keep the timber supply. Also, this buffer should be created by delaying the harvest of young and less vulnerable stands. These results are in line with Reed & Errico (1986). Cairns *et al.*, (2008) used the LANDIS model to show that a landscape structure can be influenced by the severity and extent of an insect outbreak. The landscape was artificially created and the insect studied was southern pine beetle (*Dendroctonus frontalis* Zimmermann). This has implications for the management policy, but as far as we know no follow-up study on optimization of the management policy under stochastic insect outbreaks has been performed.

#### 1.2.4 Other forest-level related problems

One aspect commonly considered in the optimization of the strategic management policy is wildlife-related and biodiversity objectives. Models incorporating these objectives are sometimes referred to as “spatial optimization” models. These types of models try to capture spatial relations between the land areas, or stands, and how they should be managed in order to optimize wildlife and environmental objectives. A large number of models have been suggested that incorporate deterministic spatial characteristics or constraints such as patch size, habitat amount thresholds, habitat connectivity, fragmentation relationships, population growth and dispersal, proximity relations such as edge effects, adjacency constraints, green up and area restrictions. For a general overview of models focusing on

optimization of deterministic characteristics of the landscape pattern, see (Hof & Haight, 2007; Baskent & Keles, 2005; Murray & Snyder, 2000). For the stochastic part, stochastic population models describing population viability, occurrence, migration, or colonization have been used. Examples include evaluation of the effect of management regimes (Fries & Lämås, 2000), the evaluation of effect of habitat management options (Akçakaya *et al.*, 2004; Liu *et al.*, 1995; Armbruster & Lande, 1993), optimization of habitat restoration (Bever, 2007), reserve selection according to growth rates (Haight & Travis, 2008; Carroll *et al.*, 2003), and reserve selection according to probabilistic presence-absence species information (Strang *et al.*, 2006; Arthur *et al.*, 2002; Polasky *et al.*, 2000). Currently, to the best of our knowledge stochastic population models have not been incorporated into the optimization of the strategic management policy.

Some work has also been done towards the incorporation of multiple types of uncertainties and disturbances. Xi *et al.*, (2008) proposed a framework for evaluating alternative management policies and restoration practices for forest restoration under spatial natural disturbances. Using a simulation model, Wanga *et al.*, (2006a; 2006b) evaluated the effect of different planting policy on landscape forest restoration after large catastrophic events. Spring & Kennedy (2005) proposed a model for optimization of commercial and ecological values of a forest. The management of the forest was thereby optimized under the risk of damage due to fire and the risk of the disappearance of an endangered species from the stands. The problem was modeled as an MDP and the forest considered consisted of four stands, the location of which was not modeled, making the model non-spatial. The probability of damage in a stand is thereby independent of the state of the neighboring stands. The same type of model was also used by Spring *et al.*, (2008) to optimize the management of a forest consisting of four stands for timber production and the maintenance of nesting sites for wildlife.

Strategic planning according to climate change and climate warming has also been given much attention recently. As climate warming influences the ecological processes on a number of scales (Iverson *et al.*, 1999; Sykes & Prentice, 1996; Ritchie, 1986), strategic planning according to unknown future climate scenarios has been studied. Bua *et al.*, (2008) used the LANDIS model to evaluate forest harvesting and planting strategies under scenarios of warmer climate. Schumacher & Bugmann (2006) evaluated the long-term effects of different climate, natural disturbances, and management strategies on the dynamics of the forest vegetation. Spring *et al.*, (2005b; 2005a) studied the management of forested catchment under changes in

rainfall runoff and in frequency of fire due to climate change. The optimal rotation age for a single even-aged stand was computed under different climate change scenarios utilizing stochastic dynamic programming. However, no spatial consideration to the location or interactions between the stands was given in these models.

### 1.2.5 State of the art

A number of models for incorporation of a variety of sources of risk and uncertainty in strategic forest planning have been proposed. However, to the best of our knowledge, there is currently no method of optimizing a large-scale management policy under stochastic events that also incorporate spatial dependencies between the stands. A number of models considering small-scale forestry, ranging from 2 to 16 stands, have been proposed. A common shortcoming of these models is, however, that either they do not consider the spatial structure of the uncertainties, or they scale exponentially with the number of stands. At the large-scale forest-level, numerous models have also been suggested. However, they either do not consider the spatial structure of the uncertainties or they can only be used to evaluate different management strategies. A number of large-scale forest-level models have been proposed in which the forest is aggregated according to age-classes, and in which the forest develops in either a deterministic or a stochastic manner. However, as the forest is aggregated according to age classes, the model cannot account for the spatial structure of the uncertainties. No spatial dependencies between stands, or varying intensity or severity of damage occurrences are considered in the models. To incorporate spatial dependencies, numerous large-scale forest-level simulation-based models have been proposed. An advantage of these models is that advanced simulation models of forest, the ecological processes, and damage occurrences can be used to simulate the future state of the forest. Thus, they can be efficiently used to evaluate different management strategies. However, to the best of our knowledge there is no simulation-based model that optimizes the management policy under stochastic development of the forest. The management policies that have been evaluated are either predefined by the user or optimized under deterministic development of the forest.

## 1.3 Markov Decision Processes

*Markov decision processes* (MDPs) (Sutton & Barto, 1998; Bertsekas & Tsitsiklis, 1996; Puterman, 1994) have been used for some time to model

and solve sequential decision-making problems under conditions of uncertainty in forestry, natural resources management, and agricultural economics (Kennedy, 1986). It has been successfully applied to numerous forest planning problems under uncertainty, ranging from climate change (Spring *et al.*, 2005b), reserve site selection (Sabbadin *et al.*, 2007), risk of forest fire (Spring & Kennedy, 2005; Garcia & Sabbadin, 2001), optimal stand management under growth and price uncertainty (Zhou *et al.*, 2008; Insley & Rollins, 2005; Rollin *et al.*, 2005; Lohmander, 2000; Lin & Buongiorno, 1998), and maintenance of wildlife (Spring *et al.*, 2008). The main strength is the ability to explicitly model uncertainty concerning the future occurrence and the effect of management, allowing a flexible optimization of management strategies. In the MDP model, the uncertainties are expressed with a transition model or probability distributions, and the objectives of the decision-maker are expressed with utility or reward functions. Typically, a management policy that optimizes the expected sum of discounted rewards is computed instead of a management policy that reaches a fixed goal with certainty.

An MDP is a model of a system or environment, for example a landscape, forest, or stand. In this section, we illustrate the MDP approach with an example of an MDP model of a stand. The state of the environment is described with a state variable, and all the possible states of the system are defined by the state space. The decision-maker, or agent, interacts with the systems by selecting and performing management activities, or actions, on the environment. Examples of actions in the stand case are clear-cutting of the stand, performing a thinning in the stand, or simply leaving the stand to grow. The action selected by the decision-maker to be performed is described by an action variable, and all the possible actions that can be performed are defined by the action space. An MDP is a sequential decision-making model, in that the decision-maker repeatedly interacts with the environment and in that the state of the environment evolves iteratively. The environment is assumed to evolve in discrete time and the interaction between the decision-maker and the environment consists of several steps. First, the decision-maker observes the current state of the environment. We will assume that the state of the system is fully observable and that the decision-maker will always observe the true state of the environment. Based on the state of the environment, the decision-maker selects and performs an action on the environment. Based on the state of the environment and the selected action, a reward is received—for example, representing the timber harvested by clear-cutting of the stand or the cost of performing a regeneration of the stand. After the action has been performed, the

environment stochastically evolves to a new state. The development of the environment is defined by a stochastic transition model, which is based on the current state of the stand and on the action taken by the agent. After the development of the environment, the agent observes the new state of the environment and again selects an action to be performed. As the system may now be in a different state, the action selected does not have to be the same. One iterative step in this process is commonly referred to as a time period. See Figure 1 for a graphical representation of the dependencies of an MDP.

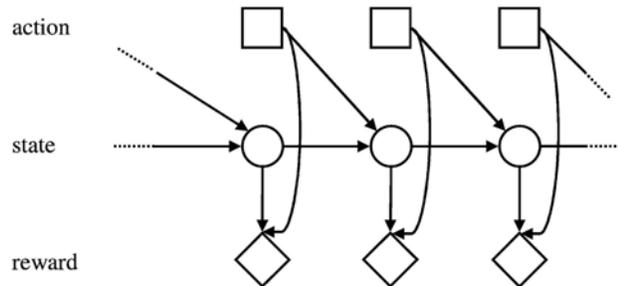


Figure 1. A Markov decision process. Square nodes represent the action, circles represent the state, and the diamond represents the reward. The reward is dependent on the state and the action, while the state is dependent on the previous state and action.

The environment is assumed to evolve in discrete time, and the aim is to define how the decision-maker should select the action to be performed, that is, to find the management policy that maximizes the expected sum of discounted rewards that will be received. A policy thus tells the decision maker what action to take for each possible state of the environment. A policy is optimal if, for each possible state of the environment, it specifies the action to be taken as the action that maximizes the expected sum of the discounted rewards that will be received. Numerous algorithms exist for computing a policy for an MDP; for an overview, we refer the reader to the review paper by van Otterlo (2005). We will differentiate between *model-free* and *model-dependent* algorithms. A model-dependent algorithm, also referred to as planning method, requires complete information concerning the transition model and reward function. They are based on the manipulation of *value functions*, a function over the states of the environment representing the expected discounted rewards that will be received from taking the optimal actions thereafter. A model-free algorithm only learns from observations of state transitions and the rewards received. There are a number of types of model-free algorithms, and we will focus mainly focus on *Q-learning* algorithms (Watkins & Dayan, 1992). A Q-learning algorithm

is based on the manipulation of a Q-function, these are action-value functions representing the expected discounted reward that will be received when taking a specific action in a state of the environment, and thereafter taking the optimal actions. An advantage of value functions and Q-functions is that if they are optimal, then the optimal policy can be greedily computed from the functions.

The optimal policy for an MDP can be computed with a number of model-dependent algorithms. The most commonly used algorithms are based on: linear programming (de Ghellinck, 1960; Manne, 1960), value iteration (Puterman, 1994; Bellman, 1957), policy iteration (Bertsekas & Tsitsiklis, 1996; Howard, 1960), and modified policy iteration (Puterman & Shin, 1978). The optimal policy can also be computed with model-free Q-learning algorithms. One of the most commonly used Q-learning algorithm is the SARSA algorithm (Watkins & Dayan, 1992). The algorithm is known to converge to the optimal Q-function when every state-action pair is sampled infinitely many times.

### 1.3.1 Factored MDP

Algorithms for computing the optimal policy for an MDP do, however, reach their limits when the state space is large. In an MDP, the state of the environment is expressed with a single state variable. This commonly becomes troublesome for large or spatial problems, as the state variable defines the complete information concerning the state of the environment. For example, in the case of a forest consisting of  $n$  number of stands and where each stand can take  $m$  different states, the state space is  $m^n$ . The state space thus grows exponentially with the number of stands that the forest consists of. There exist a number of methods to circumvent this problem, and we will focus on the commonly used *factored MDP* approach (Boutilier *et al.*, 2000; Boutilier *et al.*, 1995). Factored representation of MDPs was proposed by (Boutilier *et al.*, 1995); this is an approach in which the structure of the environment is exploited to express the model of the environment in a compact form. Instead of modeling the state of the environment with a single state variable, the structure of the environment is exploited to define a set of state variables, each expressing a part, or section, of the environment. In the case of a forest consisting of a number of stands, the state of each stand can for example be represented by a state variable. A state of the factored MDP is thus a description of the value of each state variable. The state space of the factored MDP is thus multidimensional and defined as the cross-product of a set of state variables.

The factored decomposition of the state space is followed by a similar decomposition of the transition and reward models. The transition of a state variable is usually not dependent on all state variables, only a small number of state variables. For example, the probability of a stand being damaged by wind is independent the state of a stand in another part of the forest. The dependencies among the state variables can be represented by a *dynamic Bayesian network* (DBN) (Dean & Kanazawa, 1989), according to which the transition model can be decomposed. The DBN thus represents the decomposition of the transition model, where the transition of a state variable is defined by a local transition function that is only dependent on a small number of state variables. The same type of decomposition can be performed for the reward function, which is decomposed as a sum of local reward functions, each of which is only dependent on a small number of variables.

Standard solution algorithms for MDPs are, however, of limited use for solving factored MDPs. As the size of the state space of a factored MDP is exponential in the number of state variables, the space needed to represent the value functions and the time needed to compute them are also exponential in the number of state variables. Unfortunately, there is no general guarantee that the structure of the factored MDP is reflected in the structure of the optimal value functions (Koller & Parr, 1999). Boutilier *et al.*, (1995) proposed a structured policy iteration algorithm in which a tree structure is used to represent value functions and policies. Policy iteration is used to update the policy and value functions, and accordingly to reshape the tree structure. A modification of the algorithm was proposed by Hoey *et al.*, (1999), who showed that instead of a tree structure, an *algebraic decision diagram* (ADD) (Bahar *et al.*, 1993) could be used, allowing a more compact form of representation. Feng & Hansen (2002) combined the ADD algorithm by Hoey *et al.*, (1999) with the heuristic dynamic programming algorithm LAO\* (Hansen & Zilberstein, 2001). Kim & Dean (2003) presented an algorithm in which the state space is also aggregated using a tree structure. Each aggregated block is treated as a single state and the algorithm successively decreases the aggregation of the states by block separation.

An approach for solving factored MDPs that has recently been given a lot of attention is the *approximate linear programming* (ALP) approach (Schweitzer & Seidmann, 1985). With this approach, the true value function of the problem is approximated by a linear value function, giving an approximate solution to the factored MDP. The linear value function is compactly represented as a linear combination of *basis functions*, each of which is of

smaller scope than the true value function. The manipulation of a value function has thereby been transformed into a problem of weighting different basis functions such that the linear value function yields a good approximation of the true value function of the problem. Koller & Parr (2000; 1999) proposed the use of *factored linear value functions* in which the basis functions are defined over a small number of state variables. Guestrin (2003) built on this idea and presented two algorithms based on the use of factored linear value functions. One algorithm is based on approximate linear programming and the other is based on approximate dynamic programming. However, the running time of the algorithms and quality of the computed solution are dependent on the basis functions, which are fixed and predefined by the designer. Poupart *et al.*, (2002) proposed an approach that fully automatically selects and modifies the basis functions. Schuurmans & Patrascu (2001) suggested an iterative constraint-generation approach in which constraints are iteratively added until a feasible solution can be found. de Farias & Van Roy (2004) devised a constraint-sampling approach in which only a sample of the constraints is considered in the ALP. However, efficient sampling requires a good distribution over the set of constraints. Dolgov & Durfee (2006) focused on a dual LP formulation of the problem and developed a composite-ALP approach that optimizes the primal and dual variables symmetrically. The benefit of the approach is that it simultaneously optimizes the linear value functions and the feasible region of the LP.

Numerous model-free algorithms have also been proposed for computing policies for factored MDPs. Kearns & Singh (1998) proposed a Explicit Explore or Exploit ( $E^3$ ) algorithm that achieves near-optimal performance in time polynomial in the state space of the MDP. The algorithm iteratively learns by updating the parameters of a model of the environment and by exploration and exploitation of the environment. The algorithm can give theoretical bounds on the value of computed policy. Brafman & Tennenholtz (2002) presented a simpler and more general version of the  $E^3$  algorithm called R-MAX. Kearns & Koller (1999) extended the  $E^3$  algorithm to the case where a DBN representation of the factored MDP is known. Their DBN- $M^3$  algorithm utilizes the structure of the DBN to estimate its parameters, and thereby learns a near-optimal policy. The algorithm is polynomial in the number of parameters of the DBN, and is thereby considerably faster than the  $E^3$  algorithm. Guestrin *et al.*, (2002b) extended the work on the  $E^3$  and R-MAX algorithms and presented an algorithm-directed factored reinforcement learning approach in which ALP

is used to direct exploration and exploitation of the parameters of the estimated model of the environment.

Recently, model-free algorithms combining factored MDPs and hierarchical methods have been presented. In a hierarchical method is commonly the notion of actions extending to consider sequences of actions, or policy, that may last for a number of time periods. The MDP framework is extended to *semi-markov decision processes* (SMDPs) (Parr & Russell, 1998; Mahadevan *et al.*, 1997; Bradtke & Duff, 1994), in which actions can have variable duration or consist of a sequence of actions. Three well-known algorithms are the MAXQ (Dietterich, 2000), HASSLE (Bakker & Schmidhuber, 2004), and VISA (Jonsson & Barto, 2006) algorithms. In the VISA algorithm, the structure of the factored MDP and the DBN is explored by the creation of a casual graph, a graph representing the conditions required to change the value of the state variables. The casual graph is used to select and improve the hierarchically structure the actions, and to perform state—abstraction from which the policies can be compactly represented and computed. Another algorithm that combines hierarchical structuring with factored representation of the MDP is that of Diuk *et al.*, (2006). The algorithm combines the algorithm-directed factored reinforcement learning approach presented by Guestrin *et al.*, (2002b) with hierarchical structuring of the policies to reduce the number of samples required to compute near-optimal policies.

### 1.3.2 Collaborative multiagent MDPs

By exploiting the factorization of the state space, factored MPDs can be used to model and efficiently compute optimal or approximate policies for environments with large state space. It is, however, more difficult to efficiently compute approximate policies for environments with large state and action spaces (Garcia & Sabbadin, 2001). Commonly, only the state space is factorized in a factored MDP, and most solution algorithms for factored MPDs scale poorly with the action space. Fortunately, in a number of real-world agricultural management problems, the structure of the environment can be used to factorize both the state and action spaces. The use of *collaborative multiagent MDPs* (Guestrin, 2003) is an approach that is similar to factored MDPs, but in which both the state and action spaces are factored. It is called “multiagent” as it assumes that the environment is being managed by a number of agents, or decision-makers, each observing part of the environment and each specifying what actions should be performed in their part of the environment. For example, a group of hunters working together to locate the game. Each hunter oversees a part of the forest, and at

the end of the day the whole group shares the meat. It is called “collaborative” as the agents try to work together to achieve a common goal. Depending on the environment, the agents may observe the state of each other’s part of the environment. Furthermore, the agents may communicate and coordinate concerning the selection of actions in order to maximize the overall common goal. In a collaborative multiagent MDP, the state and action spaces are multidimensional and are defined as cross-products of a set of state and action variables, respectively. A *dynamic decision network* (DDN) (Dean & Kanazawa, 1989), is used to represent the dynamics of both the transition and rewards. Also, the transition and rewards are decomposed and the local functions only depend on a small number of state and action variables.

A number of model-dependent algorithms based on the ALP approach have been suggested for computation of policies for multiagent cooperative MDPs. Guestrin *et al.*, (2001) devised a model-dependent LP approach for computing approximate policies. The number and structure of the basis functions are selected by the user, and the solution and running time of the algorithm is thereby dependent on this selection. Guestrin *et al.*, (2002b) presented an algorithm in which the factored MDP is divided into subsystems that interact in a simple manner. The systems only overlap with each other for some state variables, and a policy for each system can therefore be computed and optimized by a message passing scheme. de Farias & Van Roy (2004) proposed a version of their constraint-sampling approach for handling problems with large, or exponential, action space. However, they noted that for the algorithm to perform efficiently, it might need an even larger number of basis functions than for a factored MPD.

Regarding model-free algorithms, a number of algorithms have also been suggested. Peshkin *et al.*, (2000) suggested a direct policy search algorithm in the partially-observable setting, but which is also applicable to the complete observable setting that is assumed in this thesis. Sallans & Hinton (2004; 2000) approximated the Q-functions with a product of experts, probabilistic models that combine simpler models to represent the Q-functions in an approximate way. Claus & Boutilier (1998) presented an approach in which the local Q-functions are distributed over the agents and optimized only using local information: the locally received reward and the local value function. Schneider *et al.*, (1999) presented an approach in which the local Q-functions are distributed over the agents and optimized using locally received rewards and neighboring value functions. Guestrin *et al.*, (2002a) presented an exact and distributed algorithm for coordination of the actions between the agents; the coordination problem. As the local reward functions

are dependent on a set of action variables, the agents must coordinate their action selection to maximize the common goal. The algorithm is based on variable elimination (VE) and computes the global action that optimizes the sum of the local Q-functions. The algorithm is used in three coordinated reinforcement learning approaches: Q-learning, policy iteration, and policy search. The suggested VE algorithm does, however, have exponential complexity in the largest clique generated by the algorithm. Kok & Vlassis (2004) proposed an approximate max-plus algorithm for computing the global action that maximizes the global reward. Furthermore, they proposed a number of sparse cooperative Q-learning approaches in which the local Q-functions are distributed according to the DDN, that is, over the agents or over the edges of the DDN.

## 2 Objectives and main contribution of the thesis

The main objective of this thesis has been to develop ways of efficiently computing large scale-management policies for forest planning problems that involve risks and uncertainties associated with spatial relationships. As planning in forest management commonly considers large-scale estates, it is of particular importance to develop a framework that is scalable. In addition, due to the great variety of phenomena associated with risk in forestry, the framework should be able to represent a large variety of forestry planning problems. For the framework developed, we sought to find efficient approximate algorithms that exploit the problem structure for computation of near-optimal management policies. As a large number of simulation-based models of forestry and agricultural occurrences have already been developed, we sought to develop both model-based and model-free algorithms. Of particular importance are algorithms that are scalable and able to efficiently compute near-optimal policies for large-scale management problems. The main contribution of the thesis is as follows:

- **Modeling:** A framework is presented for compact modeling of large-scale stochastic and spatial forest planning problems. The framework of *graph-based Markov decision processes* (GMDPs) is based on MDPs and utilizes a graph structure to represent the local dependencies. The framework exploits the local dynamics of the environment to compactly express the transition and reward functions as local functions, each of which is dependent on a single action variable.
- **Efficient model-based algorithms:** We present and demonstrate the performance of two approximate algorithms that exploit the problem structure to efficiently compute near-optimal policies. The algorithms are based on the manipulation of value functions and

approximate the value function with a sum of local value functions, each of limited scope. One algorithm is based on approximate linear programming and can be implemented fully distributed. The second algorithm is based on approximate policy iteration and uses a mean-field occupation measure to approximate the process by a process with a simpler dependence structure. The algorithms only scale linearly and polynomially in the size of the problem for a fixed induced width of the graph, and can thereby solve large-scale planning problems.

- **Efficient model-free algorithms:** We present and demonstrate the properties of a set of model-free algorithms for computation of management policies for large-scale stochastic and spatial forest planning problems. The algorithms are based on the manipulation of Q-functions.
- **Multiagent coordination:** We demonstrate that coordinated action selection is not required in the GMDP framework. Even though the GMDP framework is a multiagent problem, it is not necessary for the local agents in the system to coordinate their action selection. Policies can therefore be computed for GMDPs without having to consider this generally intractable problem.
- **Collaborative multiagent MDPs:** We propose an efficient model-based algorithm for collaborative multiagent MDPs. The algorithm is based on a conversion of the collaborative multiagent MDP model into a GMDP model, for which an approximate solution to the collaborative multiagent MDPs can be efficiently computed.
- **Forest planning under risk of wind damage:** We propose a model for large-scale forest planning under the risk of wind damage that is based on the GMDP framework. We demonstrate the viability of the approach for a real-world problem. In addition, the economic gain of managing an estate according to the risk of wind damage is evaluated.

## 3 Summary of Papers

### 3.1 A framework and algorithms for the local control of spatial processes (Paper I)

In paper I, we address the problem of modeling and optimizing management policies for large-scale spatial and temporal natural resource management problems. We present a general framework for modeling of these types of problems and two algorithms for computing high-quality management policies. The framework of *graph-based Markov decision processes* (Forsell & Sabbadin, 2006; Peyrard & Sabbadin, 2006) offers a compact representation of collaborative multiagent MDPs and is specifically developed for modeling and solving large-scale spatial and temporal agricultural management problems. The state and action spaces of the GMDP are factorized and multidimensional of identical dimensions. In a GMDP, there are no local dependencies between the action variables, and the framework decomposes the transition and reward functions into local functions, each of which is dependent on a single action variable. Dependencies between the state variable are represented by a graph structure  $G = (V, E)$  in which the nodes represent the stands and edges indicate local dependencies. Each state variable is represented by a node and dependencies between the state variables are represented by directed edges between the nodes. The transition and reward functions are compactly represented by the graph structure and the GMDP representation only scales linearly with the size of the state and action space.

We present two approximate algorithms for computation of policies for GMDPs. The first algorithm is based on *approximate linear programming* (ALP) (Forsell & Sabbadin, 2006), while the second algorithm is based on *approximate policy iteration* and *mean-field approximation* (MF-API) (Peyrard &

Sabbadin, 2006). The two algorithms have been specifically developed to exploit the particular structure of the GMDP representational model and only scale linearly and polynomially with the size of the problem for a fixed, induced width of the graph. The algorithms are restricted to only considering the suboptimal set of local policies and are based on the manipulation of a value function. In both algorithms, the value function is defined as a sum of local value functions, each of limited scope. In the ALP algorithm, each local value function is dependent on a single predefined basis function and on a single state variable. A decomposition approach can thus be used to compute the value of each local value function independently of the other local value functions. The value of each value function is computed with an LP of limited size, and as the value functions are independent this can be done in parallel. In the MF-API algorithm, the local value functions are dependent on a set of state variables and on a set of local value functions. These sets are defined according to the local dependencies between the state variables expressed by the graph structure of the GMDP. As the local value functions in the MF-API algorithms have larger scopes than in the ALP algorithm, the MF-API algorithm is more complex than the ALP algorithm and gives a better approximation of the value function of a local policy.

The proposed algorithms were evaluated on two planning problems derived from real-world natural resource management problems: the management of an agricultural area in which a disease is spreading and contaminating the crop fields, and the management of a forest in which the stands may be damaged by wind. Our experimental results confirm that the algorithms can efficiently compute high-quality policies for large-scale problems; near-optimal policies were computed for a forest consisting of 196 stands and for an agricultural area consisting on 500 crop fields. For both problems, the policies computed by the algorithms were of equivalent quality, showed near-optimal performance, and outperformed naive policies such as greedy or random policies. Results are presented in Figure 2 and Figure 3. Our experimental results also show that the quality of the policies computed by the two algorithms are independent of the size, connectivity, and topology of the area being considered. Furthermore, they are complementary in that the ALP algorithm is faster while the MF-API algorithm provides a higher-quality approximation of the expected value of the computed policy. For a fixed induced width of the graph, the running time of the MF-API algorithm increases polynomially with the size of the problem while the ALP algorithm only increases linearly. On the other hand, the approximate evaluation of the value of the policy computed by

the MF-API algorithm was far more precise than that of the ALP algorithm. The approximation computed by the MF-API algorithm can be used directly as an estimation, while estimation by Monte Carlo methods (Robert & Casella, 1999) may be preferred for the policy computed by the ALP algorithm.

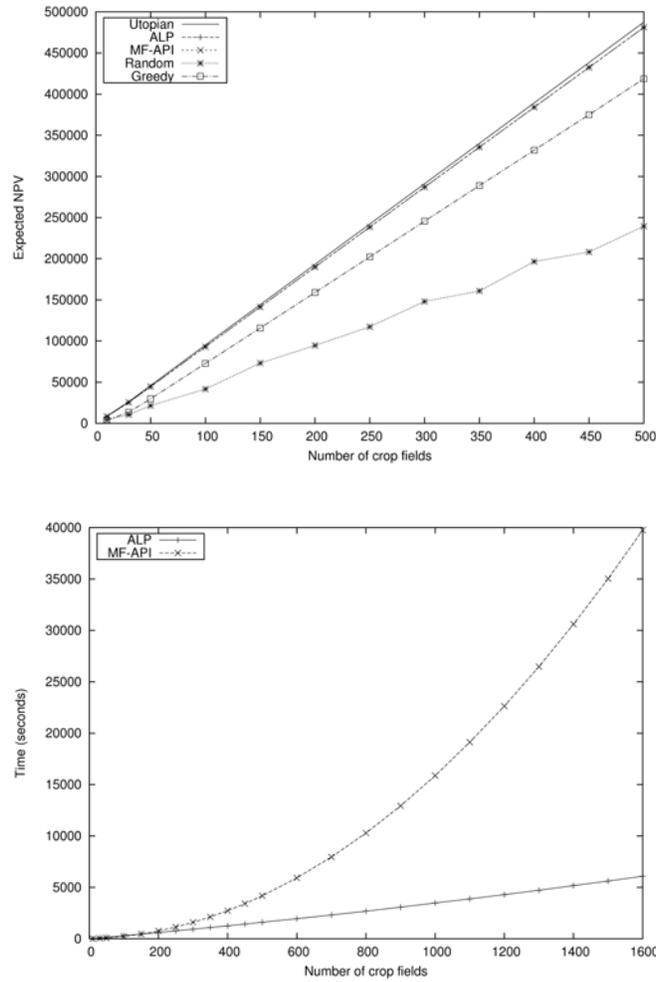


Figure 2. Evaluations of the proposed ALP and MF-API algorithms on the disease management problem. Estimated expected value of the computed policies and running times of the algorithms.

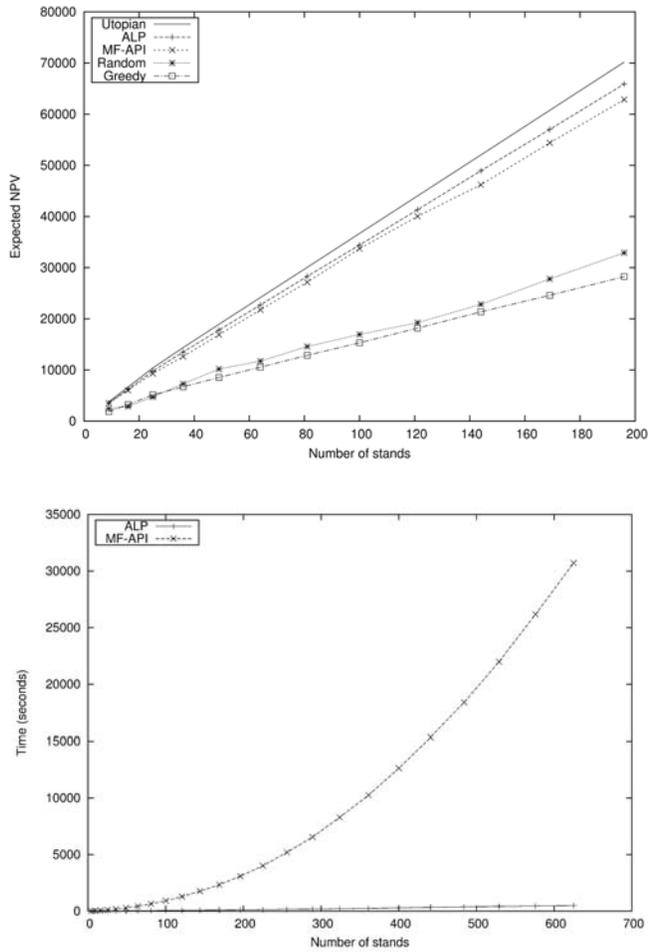


Figure 3. Evaluations of the proposed ALP and MF-API algorithms on the wind management problem. Estimated expected value of the computed policies and running times of the algorithms.

### 3.2 Q-learning for graph-based Markov decision processes (Paper II)

In paper II, we present a number of reinforcement-learning algorithms for computation of high-quality policies for GMDPs. The algorithms are model-free and do not require complete knowledge concerning the transition and reward functions. We consider RL algorithms proposed for

factored MDPs and collaborative multiagent MDPs, and show how they can be adapted to the GMDP framework. With the proposed RL algorithms, we investigate the value of coordinating the action selection, the coordination problem, within the GMDP framework. Coordinated action selection was not considered in the ALP and MF-API algorithms proposed in paper I, and as this may lead to suboptimal policies, we investigate the possible gain in quality of the resulting policy by coordinating the action selection. Furthermore, we present an efficient and scalable model-based approximate algorithm for computation of policies for a specific class of collaborative multiagent MDPs. The algorithms use the similarities between collaborative multiagent MDPs and GMDPs to convert the collaborative multiagent MDP problem into a GMDP problem. Thereafter, a policy is computed for the GMDP problem with the algorithms presented in paper I, giving an approximate solution to the collaborative multiagent MDP.

To assess the quality of the proposed RL algorithms, they were compared to the model-dependent ALP and MF-API algorithms proposed in paper I using three problems. First, there was the problem of disease management in an agricultural area consisting of eight crop fields. Second and thirdly, there were problems of forest management under risk of wind damage. In the second problem, the forest consisted of 9 stands and in the third problem the forest consisted of 100 stands. Results are presented in Figure 4. The results for the disease management problem show that with the specified number of learning iterations, the policies computed with RL algorithms had a lower estimated value than the policies computed by the model-dependent ALP and MF-API algorithms. However, the results show that several of the RL algorithms were iteratively improving the quality of their policy. For the 9-stand forest management problem, the results showed that a number of algorithms were able to compute policies that in terms of estimated value were similar to the policies computed by the model-dependent algorithms proposed in paper I. As some of the RL algorithms were not scalable to the 100-stand forest management problem, only some of the proposed RL algorithms were evaluated using the 100-stand forest management problem. Again, some of the RL algorithms were able to compute policies that in terms of estimated value were similar to the policies computed by the model-dependent ALP and MF-API algorithms. Our experimental results also showed that coordination of the action variables did not increase the expected value of the resulting policy. Two RL algorithms computed policies with and without coordinated action selection, and neither computed a policy with a higher estimated value by coordinating the action selection.

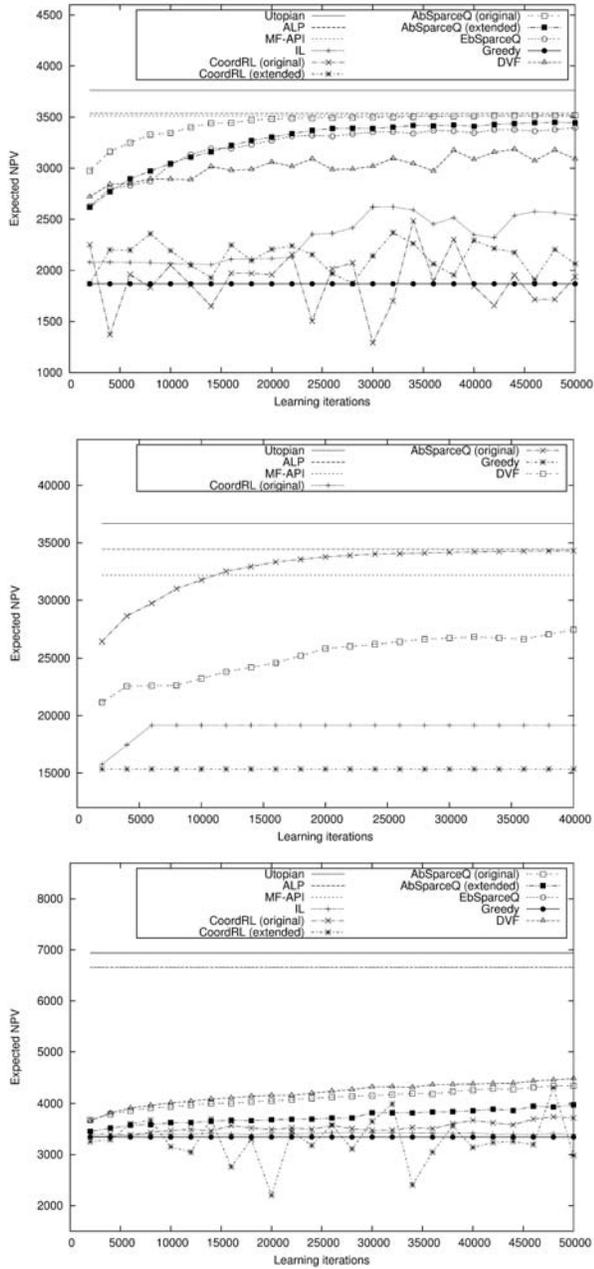


Figure 4. Performance of the RL algorithms on the problems: (top) management of a forest consisting of 9 stands under the risk of wind damage, (middle) management of a forest consisting of 100 stands under the risk of wind damage, (bottom) disease management in a agricultural area consisting of 8 crop fields.

To evaluate the proposed approximate algorithm for collaborative multiagent MDPs, we compared it to RL algorithms for collaborative multiagent MDPs. The proposed algorithm converted the collaborative multiagent MDPs into a GMDP, for which a policy was computed with the ALP and MF-API algorithms. The results are presented in Figure 5. They show that the loss in expected value of the policy imposed by the proposed conversion algorithm was approximately 20%. The highest expected value of the policies computed for the collaborative multiagent MDP was approximately 20% higher than the highest expected value of the policies computed for the GMDP.

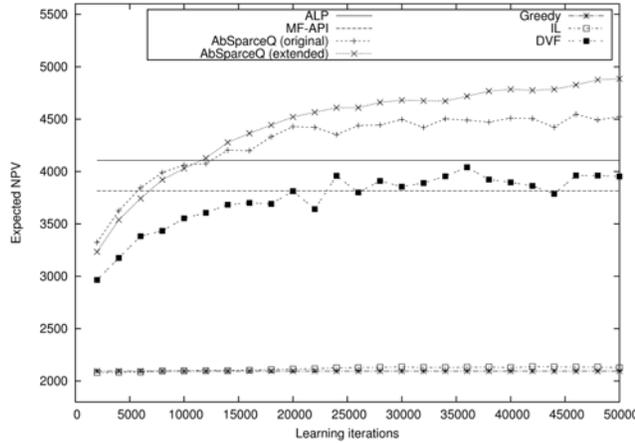


Figure 5. Evaluations using the proposed conversion algorithm (ALP, MF-API) and RL algorithms for a collaborative multiagent MDPs formulation of the forest management problems under risk of wind damage. The forest consisted of 9 stands.

### 3.3 Management of the risk of wind damage in forestry: a graph-based Markov decision process approach (Paper III)

In paper III, we present how the GMDP framework can be used to include risk, uncertainty, and spatial dependencies in long-term forest management. We propose a GMDP model for the problem of optimizing long-term silvicultural management policies under stochastic wind damage events. The model is capable of responding to stochastic wind events, the probability of wind damage, the spatial structure of the forest, the state of the neighboring stands, and the geographic orientation of the forest. As the algorithms for optimizing management policies for GMDPs only scale

linearly and polynomially with the size of the problem, the proposed model can be applied to forest estates consisting of a large number of stands. The model is demonstrated for an estate in southern Sweden covering 2,800 ha, of which 1,200 ha is forest land. The forest land is divided into 623 stands that are mainly dominated by Norway spruce. See Figure 6 for an overview of the Björnstorp estate.



Figure 6. The Björnstorp estate. The gray areas represent the forest stands and the white areas represent the non-forest areas.

The estate was modeled with twenty-year-long time periods, and for each stand two management activities could be selected: to clear-cut the stand or not to clear-cut the stand. Apart from these two management activities, the stands were also treated according to a number of fixed management activities such as site preparation, planting, pre-commercial thinning, and thinning. A growth-and-yield simulator (Wikström, 2000) was used to compute revenues and characteristics for each stand. Thus, a table was generated for each stand that specified the characteristics of the stand for each possible state of the stand, and the revenues from and costs of different activities, including the revenue if the stand was salvage-harvested. The growth-and-yield simulator used time periods of five years, the results of which were then aggregated to twenty-year periods by taking the mean over

the corresponding five-year periods. For example, the clear-cutting revenue in a certain twenty-year period would be the average of clear-cutting revenues from the corresponding five-year periods.

The proposed GMDP model considers stochastic wind events and each stand may be damaged by wind during a time period. If a stand is damaged by wind, it is assumed that the stand will always be salvage-harvested. To estimate the probability of wind damage for the stands, a tool developed by Olofsson & Blennow (2005) was used. The tool was created for classifying the annual probability of wind damage on the edges of a stand through a classification procedure based on the state of the forest and on the geographical location of the stand. Each edge section of the stand could thus be classified to have either a high or a low annual probability of being damaged by wind. Based on the classification of the edge section of the stand, the probability of the stand being damaged by wind could be computed. For simplicity, we assumed that the landscape was completely flat. Topographic shelter was thus ignored and all stands were assumed to be sheltered in terms of both large- and small-scale variations in the terrain.

Thanks to the structure of the estate and the assumptions for the tool used to evaluate the probability of wind damage, the GMDP model of the estate could be expressed in a compact form. The graph structure  $G = (V, E)$  was defined to represent – for each stand – which of the stands would have an influence on its probability of being damaged by wind. The neighboring dependencies between the stands were defined according to the borders between the stands. If two stands shared a common border, then the directed edges between the nodes representing the two stands were defined in the graph structure. However, if two neighboring stands did not influence each other's probability being damaged by wind, then the directed edges between the stands were redundant and could be removed from the graph structure. The removal of redundant and uninfluential neighbor dependencies was first performed according to the road network and non-forest areas. The tool used to classify the annual probability of wind damage is based on the assumption that if two stands are separated by a road or a non-forest area, then the stands are mutually independent of each other and cannot provide each other with shelter from wind. Secondly, neighbor dependencies between neighboring stands that were always classified as having a low annual probability of wind damage were removed. If, for all possible combinations of the state of the stand itself and its neighbor, the stand was classified as having a low annual probability of wind damage, then the state of the stand was independent of the state of the neighboring stand.

To evaluate the value of recognizing the risk of wind damage when computing a management policy, two silvicultural management policies were computed for the estate: a wind policy taking the risk of wind damage into account, and a no-wind policy ignoring the risk of wind damage. In the wind policy, the stands were treated according to their individual risk of being damaged by wind, while for the no-wind policy it was assumed that the stands would never be damaged by wind and were therefore not treated according to the risk of wind damage. As the precise probability of wind damage was unknown, wind and no-wind management policies were computed and evaluated for different probabilities of wind damage. The economic effect of managing the estate according to the wind and no-wind management policies for different probabilities of wind damage is presented in Figure 7.

Managing the estate according to the risk of wind damage increased the expected net present value (NPV) of the estate by less than 2%. Our experimental results showed that the edge sections of 60% of the stands were always classified as having a low annual probability of being damaged by wind, and that relatively few of the stands (133 of 623) were treated differently between the wind and no-wind management policies. The increase in expected NPV for the stands that were managed differently between the wind and no-wind management policies was, however, greater, from 3% to as much as 8%, depending on the probability of wind damage. Most of the stands were clear-cut earlier in the wind policy than in the no-wind policy, or during the same time period, for all risk levels evaluated. However, some stands were clear-cut later in the wind management policy, while other stands were never clear-cut.

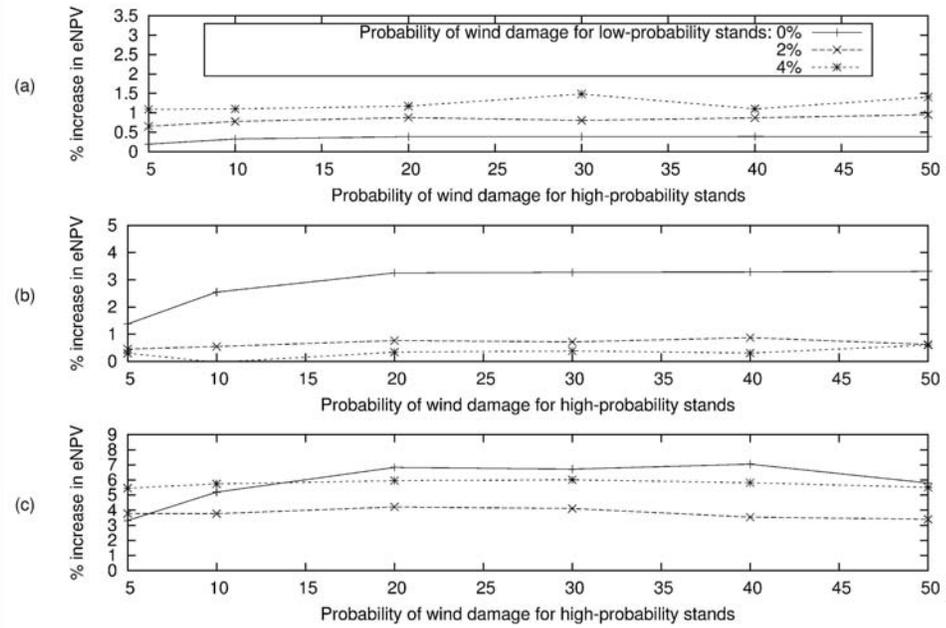


Figure 7. The percentage increase in expected net present value (NPV) by managing the estate under the wind policy rather than the no-wind policy. The top section (a) shows the increase for the whole estate, (b) shows the increase for the stands subject to the risk of damage by wind, and (c) shows the increase for stands that were treated differently in the wind and no-wind management policies.

## 4 Discussion and Conclusion

In this thesis, it has been shown that by exploiting the problem structure, we can compute high-quality management policies for large spatial resource management problems under risk and uncertainties. Interestingly, is it the specific structure of the spatial aspects of the risk and uncertainties that allow us to compactly represent and efficiently compute policies for this type of problem. We show that even though the spatial aspects add to the complexity of the problem, the local aspect of the spatial dependencies makes it possible to efficiently compute management policies.

### 4.1 Modeling

In paper I, we studied graph-based Markov decision processes, which form a special case of collaborative multiagent MDPs. The framework can be used to model large-scale spatial resource management problems, and addresses what we believe to be a common case in forestry and agricultural planning where an action only affects a single state variable. The main difference between collaborative multiagent MDPs and GMDP is that in a GMDP the local transition and reward functions are only dependent on a single action variable. The framework thus facilitates the computation of high-quality management policies for large-scale problems.

In contrast to the models proposed by Spring *et al.*, (2005b), Hof *et al.*, (1997), Boychuk & Martell (1996), Gassmann (1989), Reed & Errico (1986), Van Wagner (1983), the GMDP is stochastic and explicitly models the position and spatial dependencies between the stands. The GMDP framework is similar to the first model proposed by Meilby *et al.*, (2001). However, the model proposed by Meilby *et al.*, (2001) considers a finite number of time periods and scales exponentially with the number of stands, and polynomially or exponentially with the number of time periods. The

GMDP model, on the other hand, considers the infinite case and only scales linearly with the number of stands. Furthermore, in the model proposed by Meilby *et al.*, (2001), the local dependencies between the stands are only defined according to the borders between the stands. If two stands share a common border, then they can provide shelter for each other. Stands at a distance from each other are thus assumed not to influence each other's probability of being damaged by wind. This may be limiting in the case of small stands or rough terrain. In the GMDP model, the neighboring relations are defined according to the graph structure. The graph structure is defined by the user, and if a stand is believed to influence the evolution of a stand, an edge between the stands can simply be added. Thus, stands at a distance from each other can be defined to influence each other's development. The GMDP framework can thereby be used to express a large variety of natural resource management problems in which the aspect of non-neighboring dependencies is important to consider. One such example is forest management under risk of damage due to insect where the insects, may spread through a large area during a single time period.

## 4.2 Efficient model-based algorithms

In paper I, we studied approximate solution algorithms for GMDPs, that is, algorithms for computation of management policies for GMDP problems. As optimal policies can only be computed for small-scale problems, we present two approximate algorithms that can be used to solve large-scale problems. The ALP algorithm (Forsell & Sabbadin, 2006) is based on approximate linear programming, while the MF-API algorithm (Peyrard & Sabbadin, 2006) is based on approximate policy iteration and uses a mean-field approximation of the occupation measure of the Markovian process. The algorithms are restricted to only considering local policies, and are based on the manipulation of a sum of local value functions of restricted scope. Our experimental results show that the proposed algorithms are capable of computing near-optimal policies for large-scale problems and that the quality of the policies is not substantially degraded as the size of the problem increases. For the largest crop field management problem considered, the expected value of the policy computed by our ALP and MF-API algorithms was within 2% of the estimation of the upper bound of the value of the optimal policy. For the largest forest management problem considered, the expected value of the policy computed by our ALP and MF-API algorithms was within 6% and 10% of the estimation of the upper bound of the value of the optimal policy. The near-optimal performance of the policies supports

the use of local policies when computing policies to GMDP problems. Even though we show in paper I that it is not always the case that there exists an optimal policy that is local, local policies offer a good trade-off between complexity and generality. However, a drawback of local policies is that they are commonly difficult to optimize under global constraints acting over the entire area, because they do not consider the complete state and action spaces. Constraints such as even harvest between time periods are important to consider in large-scale forest planning, and neither of the presented algorithms is currently able to consider this type of constraints.

As a GMDP is closely related to the framework of factored MDPs and collaborative multiagent MDPs, algorithms proposed for these frameworks can also be used to compute policies for GMDPs. However, in a factored MDP only the state space is usually factorized. The algorithms proposed for factored MDP therefore mainly focus on exploiting the factorization of the state space of the problem; they reach their limits when the action space is large or factored. In collaborative multiagent MDPs, both the state and action spaces are factorized, and a number of algorithms exploiting the factorization of the state and action spaces have been proposed. The most closely related work to ours is two algorithms developed by Guestrin *et al.*, (2003). The authors proposed two algorithms for computation of approximate policies for collaborative multiagent MDPs, one based on approximate linear programming and the other is based on approximate dynamic programming. Their methods are based on the use of an approximate linear value function represented as a linear combination of basis functions. The basis functions are specified by the user and defined over a small subset of variables. However, the value of the resulting policy and the running time of the algorithms increases with the number of variables in the basis functions. Also, in the ALP algorithms the policy is computed by solving a single large-scale LP. The algorithms that we propose exploit the fact that the local transition and reward functions are only dependent on a single action variable and are thereby able to compute near-optimal policies for large-scale problems. In comparison to the ALP algorithm proposed by Guestrin *et al.*, (2003), the basis functions in our ALP algorithms are predefined over a single state variable. Also, a decomposition technique can be used to decompose the LP into a set of small-scale LPs, resulting in an approximate solution to the GMDP.

### 4.3 Efficient model-free algorithms

In paper II, we present a number of reinforcement learning algorithms for GMDPs. The RL algorithms were adopted from the framework of collaborative multiagent MDPs, and in contrast to the ALP and MF-API algorithms they are model-free and do not require knowledge concerning the transition and reward functions. For the forest management problem, some of the RL algorithms were capable of computing policies that demonstrated near-optimal performance similar to policies computed by the model-dependent ALP and MF-API algorithms. However, for the disease management problem, the RL algorithms were unable to compute policies that in terms of expected value performed similarly to the policies computed by the ALP and MF-API algorithms. The experimental result showed that a number of RL algorithms iteratively improve the quality of the computed policy, and it is our belief that if we increase the number of learning-iterations, these RL algorithms could compute policies for the disease-management problem that demonstrate near-optimal performance similar to the policies computed by the model-dependent ALP and MF-API algorithms. On the other hand, this would increase the total running time of the algorithms. One reason for the slow learning rate of the algorithms for the disease management problem might be the high connectivity of the graph representing the topology of the crop fields. As it was assumed that the disease can spread between all neighboring fields, the graph representing the local dependencies between the crop fields had higher connectivity than the graph representing the local dependencies between the forest stands. The Q-functions were therefore larger in the disease problem than in the forest management problem, and there were more values that needed to be estimated by the RL algorithms.

### 4.4 Multiagent coordination

In paper II, we address the problem of coordinated action selection within the GMDP framework. Our experimental results show that coordinated action selection does not increase the expected value of the policy computed with the RL algorithms. As coordinated action selection is generally an intractable problem and a time consuming process, this is an important observation. It also gives further evidence for restricting oneself to only considering local policies within the framework of GMDPs. We note that to decrease the running time of the RL algorithms considered, the max-plus algorithms suggested by Kok & Vlassis (2006) could have been used instead of the VE algorithm Guestrin *et al.*, (2002a). The VE algorithm does not

scale well for highly connected graphs. The max-plus algorithm instead computes an approximate solution to the coordinated action selection problem and outperforms the VE algorithm with respect to computational time for highly connected graphs. The VE algorithm was selected as it is exact and always computed the optimal actions of the agents.

#### 4.5 Collaborative multiagent MDPs

In paper II, we propose an approximate algorithm for computing policies for a specific type of collaborative multiagent MDP. We empirically evaluated the estimated value of policies computed by our conversion algorithm and policies computed for the original collaborative multiagent MDP. We believe that the approach is interesting as it scales well to large size problems, even though the loss in quality of the resulting policy might be considered high. The approximate solution to the collaborative multiagent MDPs can be computed with the ALP and MF-API algorithms, which are relatively easy to implement and only scale linearly and polynomially with the size of the problem, for fixed-induced width of the graph structure. As the approach is based on the conversion of collaborative multiagent MDPs into a GMDP, the approach may be interesting for problems with a structure that is close to the GMDP framework—that is, problems where only some of the local reward functions are dependent on a small set of state variables. As the loss induced by the conversion is dependent on the number of local reward functions that are converted, this loss induced by the conversion is likely to be smaller for problems with a structure close to that of the GMDP framework.

#### 4.6 Forest planning under risk of wind damage

In paper III, we present how the GMDP framework can be used to study the management of a forest estate faced with the risk of wind damage. Experiments from the case study confirmed the applicability of the proposed GMDP approach. Even though the expected NPV of the whole forest increased by  $\leq 2\%$  by considering the risk of wind damage in the management of the estate, the model showed that for some critical stands there was an obvious benefit of considering the risk of wind damage in the management. This indicates the importance of finding and managing some of the stands in the estate according to a management policy that considers risk of wind damage, a task which the proposed GMDP model is capable of. Unfortunately, it is difficult to find the stands in an estate for which the risk

of wind damage should be taken into account in the computation of a management policy. Not only the current risk level of the stands, but also future ones, must be considered. Our experiments also show that it is difficult to generalize the change in management of the stands when the risk of wind damage is considered. Sometimes a stand managed according to risk should be clear-cut earlier, sometimes later, and sometimes never, depending on the risk level. Also, a number of the stands should be managed according to the state of the neighboring stands. Thus, the only way in which critical stands can be identified appears to be to analyze the entire estate.

It is interesting to note that the removal of redundant and uninfluential neighbor dependencies reduced the size of the GMDP model significantly. Reduction of the size of neighborhoods reduces the solution time of the algorithms used to compute the management policies. Optimal solution algorithms instead of approximate ones could be used for some of the stands, which in turn led to computed management policies of higher quality. However, as the approximate ALP solution algorithm is capable of computing near-optimal policies for large-scale problems, the management policies could probably have been computed without the removal of uninfluential neighbor dependencies.

The results of the case study indicate that further development of the model may be required. It is somewhat surprising that the improvement of the expected NPV does not increase noticeably when the probability of wind damage increases. It is possible that this can be attributed to the period length of twenty years, which limited the possibility of managing the estate according to different risk levels. A second reason is that the difference in damage probabilities after twenty years does not differ much between risk levels due to the aggregation of the annual probabilities of wind damage to time periods of twenty years. A third reason is that we probably underestimated the value of taking the risk of wind damage into account in the case study. A forest owner has a larger set of management activities at his/her disposal than the one considered in the case study. Examples of such silvicultural treatments are site preparation, pre-commercial thinning, and thinning. The composition of the stands can, of course, also be changed. These activities can be used to further adapt the management policies to the risk levels of the stands and may thereby further increase the expected NPV of the estate.

## 5 Future Research

This work has shown that the formulation of the stochastic spatial planning problem as a GMDP in combination with suitable solution methods has practical potential in large-scale natural resource management. The applicability to a real world-sized wind-damage problem attests to this fact. Another example of practical application of the GMDP framework to spatial resource management has been presented by Peyrard *et al.*, (2007), who applied the GMDP framework to the problem of canola crop management under the risk of fungal disease. Due to the generality of the GMDP framework, we believe that the GMDP framework could have practical use in numerous large-scale natural resource management problems. A number of forestry and agricultural problems have local spatial structures that can be efficiently modeled with the framework proposed. Also, a number of forest resource management problems modeled with the MDP framework (for example Spring *et al.*, 2008; Spring & Kennedy, 2005; Spring *et al.*, 2005b) could be investigated further by modeling of the problem with the GMDP framework. A common shortcoming of these models is that the positions of or local interactions between the stands are not considered in the model, a task that the proposed GMDP model is capable of.

There are several ways in which the practical applicability of the GMDP approach could be strengthened. When planning under risk and uncertainty, it is important to accurately model the risks, uncertainties, and management options. Studies concerning the shortening of the time periods and risk evaluations according to the topography of the forest would therefore be of interest. The topography of the forest can easily be taken into account in the classification process of the stands, and would lead to more accurate modeling of the risk levels of the stands. More importantly, this would not change the model. The effect of shortening the time periods is important to study, as it increases the management options of the stands and thereby the

possibilities of managing the estate according to the risk levels. By increasing the management options of the stands, we can estimate the expected NPV of the estate more accurately. However, shorter time periods will increase the state space significantly. It is unlikely that the state space would increase to the extent that the ALP and MF-API algorithms could not be used.

The proposed wind GMDP model could also be used to investigate the management of several neighboring estates. In the case study in paper III, a single estate was modeled and it was assumed that there were no neighboring estates. The stands on the border of the estate were therefore modeled to have no neighboring forest stands outside the estate. The stands on the border of the estate were therefore modeled to have no neighboring forest stands outside the estate. If the neighboring estate would consist of forest, then the stands on each side of the border could provide shelter for each other. A GMDP model of several neighboring estates could be used to answer questions concerning the value of adaptive management according to the state and management of neighboring estates, and the value of cooperative management of neighboring forest stands.

Another extension of the application would be the inclusion of values other than financial ones such as biodiversity, recreational value, etc. In this case, it would be appropriate to divide the kinds of utilities forests yield into two categories: those that can be attributed to the state of and actions in the single stand, and those that are global in the sense that they depend on the state of or actions in a set of stands. To the latter category one can, for example, associate suitability of habitat as a function of landscape patterns and harvest volume regulation constraints. To keep within the structure of the GMDP model, utilities that are local and that can be attributed to the state of the stands could be used. One such model is the Hartman (1976) model where a biodiversity value is associated with an existence value that increases with the age of the stand.

The extent to which utilities that span the entire forest can be incorporated into the GMDP model is still, however, an open question. This issue is essentially identical to the question of incorporating global constraints into the GMDP framework. To optimize the local policies according to global constraints is a challenging task, but one that might prove worthwhile to investigate as it would improve the applicability of the GMDP framework. Global constraints such as an even harvest between time periods are important to consider in forestry and the ability of the algorithms to consider these types of constraints would be an important step forward.

As numerous simulation-based models for forestry and agricultural occurrences have been and are currently being developed, continued

development of reinforcement-learning based algorithms is also important. An interesting approach to speed up the learning rate of the methods is that of algorithms based on TD( $\lambda$ ) (Sutton & Barto, 1998). TD( $\lambda$ ) algorithms have been efficiently applied to numerous real-world planning problems, and it would be interesting to investigate how they could be adapted to the GMDP framework. Another interesting approach for computation of management policies is that of hierarchical methods (Jonsson & Barto, 2006; Bakker & Schmidhuber, 2004; Dietterich, 2000). This type of algorithm could be especially efficient for forest management, as the management commonly involves long periods of leaving the stand to grow. Actions over numerous time periods may therefore be able to express the management possibilities of the forest sufficiently. Using such an approach, it might be possible to compactly represent the management possibilities of the forest, making it possible to increase the number of management options considered.

The proposed conversion algorithm for computing approximate policies for collaborative multiagent MDPs could also be strengthened. The conversion algorithm proposed can be used to compute policies for large-scale problems, but is only valid for a specific case of collaborative multiagent MDP. A more general algorithm for converting a collaborative multiagent MDP into a GMDP would be of interest to investigate. A possible approach for converting the local transition functions for a collaborative multiagent MDP into a GMDP would be that of joint state variables in the GMDP model. State variables in the collaborative multiagent MDP that have local dependencies according to the local transition functions could be represented in the GMDP model by a single state variable, representing the state of each state variable in the collaborative multiagent MDP. The local transition function in the GMDP would thus only be dependent on a single action variable.

Even though many issues concerning GMDPs remain to be studied, several important aspects of the framework have been investigated in this thesis. Foresters and numerous other planners have to deal with considerable uncertainties when developing management strategies. As the management strategies have a significant impact on the development of the environment, it is important to continue the development of models and methods for optimization of these management strategies under different kinds of risk and uncertainty.

## References

- Akcakaya, H.R., Radeloff, V.C., Mladenoff, D.J. & He, H.S. (2004). Integrating Landscape and Metapopulation Modeling Approaches: Viability of the Sharp-Tailed Grouse in a Dynamic Landscape. *Conservation Biology* 18(2), 526-537.
- Alexander, R.R. (1964). Minimizing windfall around clear cuttings in spruce-fir forests. *Forest Science* 10(2), 130-142.
- Armbruster, P. & Lande, R. (1993). A Population Viability Analysis for African Elephant (*Loxodonta africana*): How Big Should Reserves Be? *Conservation Biology* 7(3), 602-610.
- Armstrong, G.W. (2004). Sustainability of Timber Supply Considering the Risk of Wildfire. *Forest Science* 50(5), 626-639.
- Armstrong, G.W. & Cumming, S.G. (2003). Estimating the Cost of Land Base Changes Due to Wildfire Using Shadow Prices. *Forest Science* 49(5), 719-730.
- Arthur, J.L., Haight, R., Montgomery, C.A. & Polasky, S. (2002). Analysis of the Threshold and Expected Coverage Approaches to the Probabilistic Reserve Site Selection Problem. *Environmental Modeling and Assessment* 7(2), 81-89.
- Bahar, R.I., Frohm, E.A., Gaona, C.M., Hachtel, G.D., Macii, E., Pardo, A. & Somenzi, F. (1993). Algebraic decision diagrams and their applications. In: *Proceedings of ICCAD* pp. 188-191.
- Bakker, B. & Schmidhuber, J. (2004). Hierarchical reinforcement learning with subpolicies specializing for learned subgoals. In: *Proceedings of Neural Networks and Computational Intelligence* pp. 125-130.
- Baskent, E.Z. & Keles, S. (2005). Spatial forest planning: A review. *Ecological Modelling* 188(2-4), 145-173.
- Bellman, R.E. (1957). *Dynamic Programming*. Princeton: Princeton University Press.
- Bergeron, Y., Leduc, A., Joyal, C. & Morin, H. (1995). Balsam fir mortality following the last spruce budworm outbreak in northwestern Quebec. *Canadian Journal of Forest Research* 25(8), 1375-1384.
- Bertsekas, D.P. & Tsitsiklis, J.N. (1996). *Neuro-dynamic programming*. Athena Scientific.
- Bevens, M. (2007). A chance constraint estimation approach to optimizing resource management under uncertainty. *Canadian Journal of Forest Research* 37(11), 2270-2280.

- Boutilier, C., Dearden, R. & Goldszmidt, M. (1995). Exploiting structure in policy construction. In: *Proceedings of IJCAI* pp. 1104-1111.
- Boutilier, C., Dearden, R. & Goldszmidt, M. (2000). Stochastic Dynamic Programming with Factored Representations. *Artificial Intelligence* 121(1), 49-107.
- Boychuk, D. & Martell, D.L. (1996). A multistage stochastic programming model for sustainable forest-level timber supply under risk of fire. *Forest Science* 42(1), 10-26.
- Bradtke, S.J. & Duff, M.O. (1994). Reinforcement Learning Methods for Continuous-Time Markov Decision Problems. In: *Proceedings of NIPS* pp. 393-400.
- Brafman, R.I. & Tenenbholz, M. (2002). R-MAX - A General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning. *Journal of Machine Learning Research* 3, 213-231.
- Bua, R., Hea, H.S., Hua, Y., Changa, Y. & Larsenc, D.R. (2008). Using the LANDIS model to evaluate forest harvesting and planting strategies under possible warming climates in Northeastern China. *Forest Ecology and Management* 254(3), 407-419.
- Cairns, D.M., Lafon, C.W., Waldron, J.D., Tchakerian, M., Coulson, R.N., Klepzig, K.D., Birt, A.G. & Xi, W. (2008). Simulating the reciprocal interaction of forest landscape structure and southern pine beetle herbivory using LANDIS. *Landscape Ecology* 23(4), 403-415.
- Carroll, C., Noss, R.F., Paquet, P.C. & Schumaker, N.H. (2003). Use of population viability analysis and reserve selection algorithms in regional conservation plans. *Ecological applications* 13(6), 1773-1789.
- Caulfield, J.P. (1988). A Stochastic Efficiency Approach for Determining the Economic Rotation of a Forest Stand. *Forest Science* 34(2), 441-457.
- Černý, V. (1985). Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *Journal of Optimization Theory and Applications* 45(1), 41-51.
- Church, R.L. (2007). Tactical-level Forest Management Models: Bridging between strategies and operational problems. In: *Handbook of Operations Research in Natural Resources*. Springer.
- Church, R.L., Murray, A.T. & Barber, K.H. (1994). Designing a hierarchical planning model for USDA Forest Service planning. In: *Proceedings of 6th Symposium on Systems Analysis in Forest Resources*, Pacific Grove, California, September 6-9. pp. 401-409.
- Church, R.L., Murray, A.T. & Barber, K.H. (2000). Forest planning at the tactical level. *Annals of operations research* 95, 3-18.
- Claus, C. & Boutilier, C. (1998). The Dynamics of Reinforcement Learning in Cooperative Multiagent Systems. In: *Proceedings of AAAI/IAAI* pp. 746-752.
- Davis, L.S., Johnson, K.N., Bettinger, P.S. & Howard, T.E. (2001). *Forest Management: to sustain ecological, economic, and social values, 4th Edition*. Mc Graw Hill.
- de Farias, D.P. & Van Roy, B. (2004). On Constraint Sampling in the Linear Programming Approach to Approximate Dynamic Programming. *Mathematical Methods of Operations Research* 29(3), 462-478.
- de Ghellinck, G. (1960). Les problèmes de décision séquentielle. *Cahiers du Centre d'Études de Recherche Opérationnelle* 2, 161-179.

- Dean, T. & Kanazawa, K. (1989). A model for reasoning about persistence and causation. *Computational Intelligence* 5(2), 142-150.
- Dietterich, T.G. (2000). Hierarchical Reinforcement Learning with the MAXQ Value Function Decomposition. *Journal of Artificial Intelligence Research* 13, 227-303.
- Diuk, C., Strehl, A.L. & Littman, M.L. (2006). A hierarchical approach to efficient reinforcement learning in deterministic domains. In: *Proceedings of AAMAS* pp. 313-319.
- Dolgov, D.A. & Durfee, E.H. (2006). Symmetric approximate linear programming for factored MDPs with application to constrained problems. *Annals of Mathematics and Artificial Intelligence* 47(3-4), 273-293.
- Epstein, R., Karlsson, J., Rönqvist, M. & Weintraub, A. (2007). Harvest Operational Models in Forestry. In: *Handbook of Operations Research in Natural Resources*. Springer.
- Feng, Z. & Hansen, E.A. (2002). Symbolic Heuristic Search for Factored Markov Decision Processes. In: *Proceedings of AAAI/LAAI* pp. 455-460.
- Forsell, N. & Sabbadin, R. (2006). Approximate Linear-Programming Algorithms for Graph-Based Markov Decision Processes. In: *Proceedings of ECAI*, Riva del Garda, Italy pp. 590-599.
- Fries, C. & Lämås, T. (2000). *Different Management Regimes in a Boreal Forest Landscape: Ecological and Economic Effects*. Swedish University of Agricultural Sciences, Faculty of Forestry.
- Garcia, F. & Sabbadin, R. (2001). Solving large weakly coupled Markov Decision Processes: Application to forest management. In: Ghassemi, F., et al. (Eds.) *Proceedings of International congress on modelling and simulation (MODSIM)*, Canberra pp. 1707-1712.
- Gardiner, B.A., Stacey, G.R., Belcher, R.E. & Wood, C.J. (1997). Field and wind tunnel assessments of the implications of respacing and thinning for tree stability. *Forestry* 70, 233-252.
- Gassmann, H.I. (1989). Optimal harvest of a forest in the presence of uncertainty. *Canadian Journal of Forest Research* 19(10), 1267-1274.
- González, J.R., Palahí, M. & Pukkala, T. (2005). Integrating Fire Risk Considerations in Forest Management Planning in Spain: A Landscape Level Perspective. *Landscape Ecology* 20(8), 957-970.
- González, J.R., Palahí, M., Trasobares, A. & Pukkala, T. (2006). A fire probability model for forest stands in Catalonia (north-east Spain). *Annals of Forest Science* 63(2), 169-176.
- Guestrin, C. (2003). *Planning Under Uncertainty in Complex Structured Environments*. Diss. Stanford University.
- Guestrin, C., Koller, D. & Parr, R. (2001). Multiagent Planning with Factored MDPs. In: *Proceedings of NIPS* pp. 1523-1530.
- Guestrin, C., Koller, D., Parr, R. & Venkataraman, S. (2003). Efficient Solution Algorithms for Factored MDPs. *Journal of Artificial Intelligence Research* 19, 399-468.
- Guestrin, C., Lagoudakis, M.G. & Parr, R. (2002a). Coordinated Reinforcement Learning. In: *Proceedings of ICML* pp. 227-234.

- Guestrin, C., Patrascu, R. & Schuurmans, D. (2002b). Algorithm-Directed Exploration for Model-Based Reinforcement Learning in Factored MDPs. In: *Proceedings of ICML* pp. 235-242.
- Gustafson, E.J., Shifley, S.R., Mladenoff, D.J., Nimerfro, K.K. & He, H.S. (2000). Spatial simulation of forest succession and timber harvesting using LANDIS. *Canadian Journal of Forest Research* 30(1), 32-43.
- Gustafson, E.J., Zollner, P.A., Sturtevant, B.R., He, H.S. & Mladenoff, D.J. (2004). Influence of forest management alternatives and land type on susceptibility to fire in northern Wisconsin, USA. *Landscape Ecology* 19(3), 327-341.
- Haight, R. & Travis, L.E. (2008). Reserve Design to Maximize Species Persistence. *Environmental Modeling and Assessment* 13(2), 243-253.
- Hamel, D.R. & Shade, C.I. (1987). Pesticide use in forest management. *U.S. Dept. of Agriculture, Forest Service* FS-404, 11p.
- Hansen, E.A. & Zilberstein, S. (2001). LAO\*: A heuristic search algorithm that finds solutions with loops. *Artificial Intelligence* 129(1-2), 35-62.
- Hartman, R. (1976). The harvesting decision when a standing forest has value. *Economic Inquiry* 14(1), 52-58.
- Hennigar, C.R., MacLean, D.A., Porter, K.B. & Quiring, D.T. (2007). Optimized harvest planning under alternative foliage-protection scenarios to reduce volume losses to spruce budworm. *Canadian Journal of Forest Research* 37(9), 1755-1769.
- Hoey, J., St-Aubin, R., Hu, A.J. & Boutilier, C. (1999). SPUDD: Stochastic Planning using Decision Diagrams. In: *Proceedings of UAI* pp. 279-288.
- Hof, J., Bevers, M. & Kent, B. (1997). An Optimization Approach to Area-Based Forest Pest Management Over Time and Space. *Forest Science* 43(1), 121-128.
- Hof, J. & Haight, R. (2007). Optimization of forest wildlife objectives. In: *Handbook of Operations Research in Natural Resources*. Springer.
- Howard, R.A. (1960). *Dynamic Programming and Markov Processes*. Cambridge: MIT Press.
- Insley, M. & Rollins, K. (2005). On solving the multirotational timber harvesting problem with stochastic prices: a linear complementary formulation. *American Journal of Agricultural Economics* 87(3), 735-755.
- Iverson, L.R., Prasad, A. & Schwartz, M.W. (1999). Modeling potential future individual tree-species distributions in the eastern United States under a climate change scenario: a case study with *Pinus virginiana*. *Ecological Modelling* 115(1), 77-93.
- Johnson, K.N., Sessions, J., Franklin, J. & Gabriel, J. (1998). Integrating Wildfire into Strategic Planning for Sierra Nevada Forests. *Journal of Forestry* 96(1), 42-49.
- Jonsson, A. & Barto, A.G. (2006). Causal Graph Based Decomposition of Factored MDPs. *Journal of Machine Learning Research* 7, 2259-2301.
- Kearns, M.J. & Koller, D. (1999). Efficient Reinforcement Learning in Factored MDPs. In: *Proceedings of IJCAI* pp. 740-747.
- Kearns, M.J. & Singh, S.P. (1998). Near-Optimal Reinforcement Learning in Polynomial Time. In: *Proceedings of ICML* pp. 260-268.
- Kennedy, J.O.S. (1986). *Dynamic Programming - Applications to Agriculture and Natural Resources*. Elsevier Applied Science Publishers.
- Kim, K.-E. & Dean, T. (2003). Solving factored MDPs using non-homogeneous partitions. *Artificial Intelligence* 147(1-2), 225-251.

- Kirkpatrick, S., Gelatt, C.D., Jr. & Vecchi, M.P. (1983). Optimization by Simulated Annealing. *Science* 220(4598), 671-680.
- Kok, J.R. & Vlassis, N.A. (2004). Sparse cooperative Q-learning. In: *Proceedings of ICML*.
- Kok, J.R. & Vlassis, N.A. (2006). Collaborative Multiagent Reinforcement Learning by Payoff Propagation. *Journal of Machine Learning Research*(7), 1789-1828.
- Koller, D. & Parr, R. (1999). Computing Factored Value Functions for Policies in Structured MDPs. In: *Proceedings of IJCAI* pp. 1332-1339.
- Koller, D. & Parr, R. (2000). Policy Iteration for Factored MDPs. In: *Proceedings of UAI* pp. 326-334.
- Lee, B.S., Alexander, M.E., Hawkes, B.C., Lynham, T.J., Stocks, B.J. & Englefield, P. (2002). Information systems in support of wildland fire management decision making in Canada. *Computers and Electronics in Agriculture* 37(1-3), 185-198.
- Lin, C. & Buongiorno, J. (1998). Tree diversity, landscape diversity, and economics of maple-birch forests: implications of Markov models. *Management Science* 44(10), 1351-1366.
- Liu, J., Dunning, J.B. & Pulliam, H.R. (1995). Potential Effects of a Forest Management Plan on Bachman's Sparrows (*Aimophila aestivalis*): Linking a Spatially Explicit Model with GIS. *Conservation Biology* 9(1), 62-75.
- Lohmander, P. (2000). Optimal sequential forestry decisions under risk. *Annals of operations research* 95(1-4), 217-228.
- Lohmander, P. & Helles, F. (1987). Windthrow probability as a function of stand characteristics and shelter. *Scandinavian Journal of Forest Research* 2, 227-238.
- MacLean, D.A., Erdle, T.A., MacKinnon, W.E., Porter, K.B., Beaton, K.P., Cormier, G., Morehouse, S. & Budd, M. (2001). The Spruce Budworm Decision Support System: forest protection planning to sustain long-term wood supply. *Canadian Journal of Forest Research* 31(10), 1742-1757.
- MacLean, D.A., Porter, K.B., MacKinnon, W.E. & Beaton, K.P. (2000). Spruce budworm decision support system: lessons learned in development and implementation. *Computers and Electronics in Agriculture* 27(1-3), 293-314.
- Mahadevan, S., Marchallick, N., Das, T.K. & Gosavi, A. (1997). Self-improving factory simulation using continuous-time average-reward reinforcement learning. In: *Proceedings of Fourteenth International Conference on Machine Learning* pp. 202-210.
- Manne, A.S. (1960). Linear programming and sequential decisions. *Management Science* 6(3), 259-267.
- Martell, D.L. (1980). The optimal rotation of a flammable forest stand. *Canadian Journal of Forest Research* 10(1), 30-34.
- McCarthy, M.A. & Burgman, M.A. (1995). Coping with uncertainty in forest wildlife planning. *Forest Ecology and Management* 74(1-3), 23-36.
- Mehta, S., Frelich, L.E., Jones, M.T. & Manolis, J. (2004). Examining the effects of alternative management strategies on landscape-scale forest patterns in northeastern Minnesota using LANDIS. *Ecological Modelling* 180(1), 73-87.
- Meilby, H., Strange, N. & Thorsen, B.J. (2001). Optimal spatial harvest planning under risk of windthrow. *Forest Ecology and Management* 149, 15-31.

- Meilby, H., Thorsen, B.J. & Strange, N. (2003). *Optimal spatial harvest planning under risk of windthrow*. Recent accomplishments in applied forest economics research, Kluwer Academic Publishers.
- Mladenoff, D.J. (2004). LANDIS and forest landscape models. *Ecological Modelling* 180, 7-19.
- Mladenoff, D.J., Host, G.E., Boeder, J. & Crow, T.R. (1996). *LANDIS: a spatial model of forest landscape disturbance, succession, and management*. Fort Collins, CO, USA: GIS World Books. (GIS and environmental modeling)
- Mohren, G.M.J. (2003). Large-scale scenario analysis in forest ecology and forest management. *Forest Policy and Economics* 5(2), 103-110.
- Moll, R.H.H. & Chinneck, J.W. (1992). Modeling regeneration and pest control alternatives for a forest system in the presence of fire risk. *Natural Resource Modeling* 6(1), 23-49.
- Murray, A.T. & Snyder, S. (2000). Introduction to Spatial Modeling in Forest Management and Natural Resource Planning. *Forest Science* 46(2), 153-156.
- Nealis, V.G. & Régnière, J. (2004). Insect-host relationships influencing disturbance by the spruce budworm in a boreal mixedwood forest. *Canadian Journal of Forest Research* 34(9), 1870-1882.
- Olofsson, E. & Blennow, K. (2005). Decision support for identifying spruce forest stand edges with high probability of wind damage. *Forest Ecology and Management* 207, 87-98.
- Parr, R. & Russell, S. (1998). Reinforcement learning with hierarchies of machines. In: *Proceedings of Advances in Neural Information Processing Systems 10* pp. 1043-1049.
- Pellikka, P. & Järvenpää, E. (2003). Forest stand characteristics and wind and snow induced forest damage in boreal forest. In: *Proceedings of International Conference on Wind Effects on Trees* pp. 269-276.
- Peltola, H. (1996). Model computations on wind flow and turning moment by wind for Scots pines along the margins of clear-cut areas. *Canadian Journal of Forest Research* 83(3), 203-215.
- Peltola, H., Kellomäki, S., Väisänen, H. & Ikonen, V.P. (1999). A mechanistic model for assessing the risk of wind and snow damage to single trees and stands of Scots pine, Norway spruce, and birch. *Canadian Journal of Forest Research* 29(6), 647-661.
- Persson, P. (1975). Windthrow in forests - its causes and the effect of forestry measures. *Royal College of Forestry, Department of Forest Yield Research, Stockholm, Research Notes* 36
- Peshkin, L., Kim, K.-E., Meuleau, N. & Kaelbling, L.P. (2000). Learning to Cooperate via Policy Search. In: *Proceedings of UAI* pp. 489-496.
- Peter, B. & Nelson, J. (2005). Estimating harvest schedules and profitability under the risk of fire disturbance. *Canadian Journal of Forest Research* 35(6), 1378-1388.
- Peyrard, N. & Sabbadin, R. (2006). Mean Field Approximation of the Policy Iteration Algorithm for Graph-Based Markov Decision Processes. In: *Proceedings of ECAI* pp. 595-599.
- Peyrard, N., Sabbadin, R., Pelzer, E.L. & Aubertot, J.N. (2007). A graph-based Markov decision process framework for optimising integrated management of diseases in agriculture. In: *Proceedings of MODSIM*, Christchurch, New-Zealand pp. 2175-2181.

- Polasky, S., Camm, J.D., Solow, A.R., Csuti, B., White, D. & Ding, R. (2000). Choosing reserve networks with incomplete species information. *Biological Conservation* 94(1), 1-10.
- Poupart, P., Boutilier, C., Patrascu, R. & Schuurmans, D. (2002). Piecewise Linear Value Function Approximation for Factored MDPs. In: *Proceedings of AAAI/IAAI* pp. 292-299.
- Puterman, M.L. (1994). *Markov Decision Processes*. New York: John Wiley and Sons.
- Puterman, M.L. & Shin, M.C. (1978). Modified Policy Iteration Algorithms for Discounted Markov Decision Problems. *Management Science* 24(11), 1127-1137.
- Quine, C., Coutts, M., Gardiner, B.A. & Pyatt, G. (1995). *Forest and wind: Management to minimise damage*. London: Bulletin 114, HMSO.
- Reed, W.J. (1984). The effects of the risk of fire on the optimal rotation of a forest. *Journal of Environmental Economics and Management* 11(2), 180-190.
- Reed, W.J. (1987). Protecting a forest against fire: optimal protection patterns and harvest policies. *Natural Resource Modeling* 2(1), 23-53.
- Reed, W.J. & Errico, D. (1985). Assessing the long-run yield of a forest stand subject to the risk of fire. *Canadian Journal of Forest Research* 15(4), 680-687.
- Reed, W.J. & Errico, D. (1986). Optimal harvest scheduling at the forest level in the presence of the risk of fire. *Canadian Journal of Forest Research* 16(2), 266-278.
- Reed, W.J. & Errico, D. (1987). Techniques for assessing the effects of pest hazards on long-run timber supply. *Canadian Journal of Forest Research* 17(11), 1455-1465.
- Ritchie, J.C. (1986). Climate change and vegetation response. *Plant Ecology* 67(2), 65-74.
- Robert, C.P. & Casella, G. (1999). *Monte Carlo Statistical Methods*. New York: Springer-Verlag.
- Rollin, F., Buongiorno, J., Zhou, M. & Peyron, J.L. (2005). Management of mixed species, uneven-aged forests in the French Jura: from stochastic growth and price models to decision tables. *Forest Science* 51(1), 64-75.
- Ryu, S.-R., Chena, J., Zhenga, D. & Lacroixa, J.J. (2007). Relating surface fire spread to landscape structure: An application of FARSITE in a managed forest landscape. *Landscape and Urban Planning* 83(4), 275-283.
- Sabbadin, R., Spring, D.A. & Rabier, C.E. (2007). Dynamic reserve site selection under contagion risk of deforestation. *Ecological Modelling* 201, 75-81.
- Sallans, B. & Hinton, G.E. (2000). Using Free Energies to Represent Q-values in a Multiagent Reinforcement Learning Task. In: *Proceedings of NIPS* pp. 1075-1081.
- Sallans, B. & Hinton, G.E. (2004). Reinforcement Learning with Factored States and Actions. *Journal of Machine Learning Research* 5, 1063-1088.
- Schelhaas, M., Nabuurs, G. & Schuck, A. (2003). Natural disturbances in the European forests in the 19th and 20th centuries. *Global Change Biology* 9(11), 1620-1633.
- Scheller, R.M., Domingo, J.B., Sturtevant, B.R., Williams, J.S., Rudy, A., Gustafson, E.J. & Mladenoff, D.J. (2007). Design, development, and application of LANDIS-II, a spatial landscape simulation model with flexible temporal and spatial resolution. *Ecological Modelling* 201, 409-419.
- Schneider, J.G., Wong, W.K., Moore, A.W. & Riedmiller, M.A. (1999). Distributed Value Functions. In: *Proceedings of ICML* pp. 371-378.

- Schumacher, S. & Bugmann, H. (2006). The relative importance of climatic effects, wildfires and management for future forest landscape dynamics in the Swiss Alps. *Global Change Biology* 12(8), 1435-1450.
- Schuermans, D. & Patrascu, R. (2001). Direct value-approximation for factored MDPs. In: *Proceedings of NIPS* pp. 1579-1586.
- Schweitzer, P. & Seidmann, A. (1985). Generalized polynomial approximations in Markovian decision processes. *Journal of Mathematical Analysis and Applications* 110, 568-582.
- Shifley, S.R., Thompson, F.R., Larsenb, D.R. & Dijaka, W.D. (2000). Modeling forest landscape change in the Missouri Ozarks under alternative management practices. *Computers and Electronics in Agriculture* 27(1-3), 7-24.
- Spring, D.A. & Kennedy, J.O.S. (2005). Existence value and optimal timber-wildlife management in a flammable multistand forest. *Ecological Economics* 55(3), 365-379.
- Spring, D.A., Kennedy, J.O.S., Lindenmayer, D.B., McCarthy, M.A. & Nally, R.M. (2008). Optimal management of a flammable multi-stand forest for timber production and maintenance of nesting sites for wildlife. *Forest Ecology and Management* 255(11), 3857-3865.
- Spring, D.A., Kennedy, J.O.S. & Nally, R.M. (2005a). Optimal management of a forested catchment providing timber and carbon sequestration benefits: Climate change effects. *Australian Journal of Agricultural and Resource Economics* 49(3), 303-320.
- Spring, D.A., Kennedy, J.O.S. & Nally, R.M. (2005b). Optimal management of a forested catchment providing timber and carbon sequestration benefits: Climate change effects. *Global Environmental Change* 15(3), 281-292.
- Strang, N., Thorsen, B.J. & Bladt, J. (2006). Optimal reserve selection in a dynamic world. *Biological Conservation* 131(1), 33-41.
- Sutton, R.S. & Barto, A.G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press
- Sykes, M.T. & Prentice, I.C. (1996). Climate change, tree species distributions and forest dynamics: A case study in the mixed conifer/northern hardwoods zone of northern Europe. *Climatic Change* 34(2), 161-177.
- Thompson, W.A., Vertinsky, I., Schreier, H. & Blackwell, B.A. (2000). Using forest fire hazard modelling in multiple use forest management planning. *Forest Ecology and Management* 134(1-3), 163-176.
- UNECE/FAO (2000). Forest Products Annual Market Review 2000-2001. *United Nations Economic Commission for Europe (Geneva), Food and Agriculture Organization of the United Nations (Roma)*
- Valinger, E. & Pettersson, N. (1996). Wind and snow damage in a thinning and fertilization experiment in *Picea abies* in southern Sweden. *Forestry* 69, 25-33.
- van Otterlo, M. (2005). A survey of reinforcement learning in relational domains. *Technical Report TR-CTIT-05-31, University of Twente*
- Van Wagner, C.E. (1983). Simulating the effect of forest fire on long-term annual timber supply. *Canadian Journal of Forest Research* 13(3), 451-457.
- Wanga, X., He, H.S., Li, X., Chang, Y., Hu, Y., Xu, C., Bu, R. & Xie, F. (2006a). Simulating the effects of reforestation on a large catastrophic fire burned landscape in Northeastern China. *Forest Ecology and Management* 225(1-3), 82-93.

- Wanga, X., He, H.S., Li, X. & Hu, Y. (2006b). Assessing the cumulative effects of postfire management on forest landscape dynamics in northeastern China. *Canadian Journal of Forest Research* 36, 1992-2002.
- Watkins, C.J.C.H. & Dayan, P. (1992). Technical Note Q-Learning. *Machine Learning* 8(3-4), 279-292.
- Vélez, R. (2002). Causes of forest fires in the Mediterranean basin. In: *Proceedings of EFI* pp. 35-42.
- Wikström, P. (2000). A solution method for uneven-aged management applied to Norway spruce. *Forest Science* 46, 452-462.
- Xi, W., Coulson, R.N., D., W.J., Tchakerian, M., Lafon, C.W., Cairns, D.M., Birt, A.G. & Klepzig, K.D. (2008). Landscape Modeling for Forest Restoration Planning and Assessment: Lessons from the Southern Appalachian Mountains. *Journal of Forestry* 106(4), 191-197.
- Zeng, H., Pukkala, T. & Peltola, H. (2007). The use of heuristic optimization in risk management of wind damage in forest planning. *Forest Ecology and Management* 241, 189-199.
- Zhou, M., Liang, J. & Buongiorno, J. (2008). Adaptive versus fixed policies for economic or ecological objectives in forest management. *Forest Ecology and Management* 254, 178-187.

## Acknowledgements

This thesis would not have been possible without the encouragement, support, and help of a large number of people: my supervisors, colleges at SLU and INRA, friends and family. First of all, I want to thank my four supervisors: Ljusk Ola Eriksson, Peder Wikström, Regis Sabbadin, and Frédérick Garcia. I am deeply grateful for all your guidance and for all the discussions that we have had over these years. You have all been a constant source of encouragement and have given me advice on all aspects of PhD life and studies. Thank you!

I thank my parents, Hally and Göte Forsell, who have given me endless support and encouragement. Also, my brother Magnus Forsell for being all a brother can be, and more. Emilie-Anne Guerch, for always supporting and believing in me. Your love is overpowering and I look forward to a long life together with you. Many thanks to Jean-Luc and Michèle Guerch for accepting me into your family with open arms, and for all your help with the French translation. One could not wish for better parents-in-law!

I thank all my friends and colleagues at INRA and SLU. Working together with all of you has been an inspiration. You have all helped to make my PhD studies a wonderful time and to brighten my days at the lab. I thank all the PhD-students that I meet during these years, I will definitely miss the good times we had. My deepest thanks go out to: Faure Pascale and Jackie Feau, for helping me through French administration and bureaucracy. Olivier Crespo, for all your help when I first moved to Toulouse, for introducing me to French food, and for introducing me to Emilie-Anne Guerch. Kim-Anh Lê Cao for showing me how to enjoy life in Toulouse. Erika Sallet and Arnaud Bellec, for explaining French behavior to me. Sebastien Carrere and Sonia Vautrin, for showing me high-class French culture and humor. Romain Fremez, for always speaking to me in easy and understandable French and for teaching the all-too-important and absolutely

necessary French phrases. Melanie Derre, for forcing me to speak French and for showing me how much of the language I have actually learned. Gerald Salin and Nathalie Miraglio, for your companionship and for inviting me into your world.

Finally, I wish to thank Matthias Zytnicki, Simon Boitard, and Olivier Crespo, who introducing me to M&Ms, and who showed me the benefits of always having some in the office.