A Graph-based Markov Decision Process approach for managing forests under risk of wind damage

N. Forsell[1,2], P. Wikström[1], F. Garcia[2], R. Sabbadin[2], K. Blennow[3], L. O. Eriksson[1]

[1]Department of Forest Resource Management, SLU, Umeå, SWEDEN
[2] Unité de Biométrie et Intelligence Artificielle, INRA de Toulouse, FRANCE
[3] Southern Swedish Forest Research Centre, SLU, Alnarp, SWEDEN
E-mail: nicklas.forsell@srh.slu.se

Keywords: *Forest management, wind damage risk, spatial processes, markov decision processes*

EXTENDED ABSTRACT

When selecting a long-term silvicultural management policy for a forest, a number of stochastic events can be taken into account in order to improve the management of the forest and thereby increase its value. A stochastic event that has recently been given a lot of attention to is the damage to forests due to windstorms. This is an important event to take into consideration as recent studies show that damage due to storms may seriously influence the economic value of the forest and that the risk of wind damage is influenced by the silvicultural treatments.

When the risk of wind damage is taken into account in the optimisation of the forest management policy two major problems arise. First, the ensuing management problem becomes spatial as a stand's risk of being damaged due to windstorms depends on the state of the specific stand as well as the state of its neighbouring stands. Second, for real-world applications, the problem size may become intractably large as forests containing hundreds of stands may have to be considered.

To overcome these two difficulties, we propose the use of the newly developed framework of Graph-Based Markov Decision Processes (GMDP). With this framework, the problem can both be modelled and approximately solved by an optimisation of the forest management policy. The optimisation is such that the long term economic value of the forest (expected net present value) is maximised, under the constraint that the management policy is legal according to a number of regulations in the Swedish Forestry Act. The output of the GMDP model is a state-dependent management policy that specifies for each stand which silvicultural treatment should be applied to the stand. Advantages of the proposed model are that it can both take stochastic wind and growth events into account. The risk of wind damage can also be modelled realistically to be dependent on aspects such as: the state of the stand; the state of the neighbouring stands; the spatial structure of the forest; the geographical orientation of the forest; and the topography of the area. Furthermore, as recent development in GMDP solution methods has produced algorithms that only grow linearly with the number of stands, it is likely that the proposed method can successfully be applied to forest areas containing a large number of stands.

The main objective of this article is to present a GMDP approach for optimising a forest management policy when the risk of wind damage is taken into account. We will describe the theoretical aspects of the framework and how a model of this problem can be built. A second objective of this article is to analyse the value of taking the risk of wind damage into account in the selection of the forest management policy. This by comparison of the value of a forest when it is managed according to policies taking or not taking the risk of wind damage into account.

For a case study, a forest estate in the southern part of Sweden was selected. This is a work in progress and the work has so far been concentrated on a section of the estate. The selected section consisted of 39 forest stands to a total of 99 ha. Sub-modules for projection of the state of the forest developed for the Swedish forestry were used together with a model for estimation of the probability of wind damage. The GMDP model showed which stands were at risk of being damaged, their individual risk levels and what silvicultural treatments should be applied to maximise the expected NPV of the forest. Treating the stands according to the management policy specified by the GMDP model increased the expected NPV (10^3 SEK) of the whole forest only slightly by 2% from 1303 to 1323 However, for the stands at risk of wind damage, the expected NPV increased by as much as 19% from 104 to 124.

1 INTRODUCTION

In Europe and in Sweden, wind is a major concern in forestry due to the massive amount of damage it inflicts. For example, a storm in 2005 felled approximately 75 million m^3 of softwood in the southern part of Sweden (Sondell (2006)). In Sweden alone wind and snow damage on average 4 million m^3 of wood per year (Valinger and Fridman (1997)). In all of Europe windstorms damage 18.5 million m^3 of wood per year, (Schelhaas et al. (2003)). The economic effect of wind damage may be severe as it not only leads to timber losses, but also to departures from the forest management plan. This as damage in a forest area may force the forest owner to perform unscheduled and costly thinnings as well as clear-cuttings. These activities are performed to salvage damaged areas and prevent insect attacks on broken and uprooted trees spreading to healthy trees (Schroeder and Eidmann (1993)).

Several studies have shown that the probability of wind damage is influenced by the silvicultural treatments (e.g. Persson (1975), Lohmander and Helles (1987), Quine et al. (1995), Valinger and Pettersson (1996)). Wind damage is for example most likely to occur in stands that have recently been heavily thinned or that are located next to newly created clear-cut areas (Gardiner et al. (1997), Peltola et al. (1999)). This means that the risk level depends on the forest management plan. It has also been shown that it is feasible to take into account the probability of wind damage in the optimisation of the forest management plan. This may decrease the overall risk level and increase the expected net present value (NPV) of the forest. Meilby et al. (2001), optimise a management plan while taking into account the spatial interaction between the stands as well as the geographical structure and orientation of the forest. The risk of wind damage is explicitly modelled, based on an empirical model (Lohmander and Helles (1987)) that takes into account the sheltering effect of neighboring forest stands. However, as the model explicitly states all possible future scenarios of the forest, it grows exponentially in the number of stands and polynomially in the number of time periods. Their model is therefore unsuited for forests containing a large number of stands. Another model has been suggested by (Zeng et al. (2007)), in which the problem is modelled as an unconstrained optimisation problem and where the decisions are which forest management plan to use during three 10-year periods. The objective of the model is to minimise the number of vulnerable edges while keeping a high timber harvest and an even timber flow over the planning horizon. To evaluate the effects of a forest plan, a forest growth model (SIMA, Kellomäki and Väisänen (1997)), a mechanistic wind damage model (HWIND, Peltola et al. (1999)) and a GIS software (ArcGis) are used. An approximately optimal management plan is found with the help of heuristic optimisation methods such as simulated annealing, tabu search and genetic algorithms (Heinonen and Pukkala (2004), Pukkala and Kurttila (2005)). The method is evaluated on a 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 decreased while still satisfying economic objectives. However, their model is deterministic as it neither takes wind events and stochastic development of the stand into account. It can therefore not answer the important question of how shall the forest be managed if it is damaged?

We propose the use of a Graph-based Markov Decision Process model for the optimisation of the management policy of a forest when the risk of wind damage is taken into account. Note that the proposed GMDP model computes a forest management policy and not a forest management plan. In its classical definition a management plan is time-dependent and specifies for each time period what silvicultural treatment shall be performed. A management policy is state-dependent and specifies for each possible state of the forest what silvicultural treatment shall be performed. Management policies can therefore take into account stochastic events and thereby answer not only questions concerning the prevention of wind damage, but also what shall be done after wind damage has occurred.

This article is structured as follows. First, we describe the GMDP framework. Next we present the GMDP model of the forest management problem. Then, we show how two sub-modules are used to specify forest growth, yield from silvicultural treatments and probability of storm damage. Finally, we describe an area in the south of Sweden for which the model has been evaluated and discuss the conclusions that can be drawn from this preliminary study.

2 GRAPH-BASED MARKOV DECISION PROCESSES

In this section we formalise the problem addressed in this article within a new framework. This particular framework was selected as it provides means of modelling and solving the forest problem at hand. We start by describing the theoretical framework which is used and continue with a description of the solution algorithms that have been suggested for solving a problem expressed in this model.

In its classical formulation (Puterman (1994)), a stationary *Markov Decision Process* (MDP) is a system defined by a four-tuple $< \mathcal{X}, \mathcal{A}, p, r >$ where: $\mathcal{X}$ is a finite set representing the admissible states of the system; $\mathcal{A}$ is a finite set representing the applicable actions; p is a transition function $p : \mathcal{X} \times \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ such that $p(x'|x, a)$ represents the probability of moving from state $x \in \mathcal{X}$ to state $x' \in \mathcal{X}$ when action $a \in \mathcal{A}$ is applied; and r is an "immediate" reward function $r : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ such that $r(x, a)$ represents the reward obtained in state x when action a is applied. The MDP is defined so that at each decision epoch, an action a is selected to be applied, after which the system changes from state x to state x'. The set of decision epochs is denoted by T, and from now on it will be assumed that T is infinite, $T = \{1, 2, \dots\}$. The time between two decision epochs is known as a time period. For any decision epoch t, the state x' of the process at the following decision epoch $t + 1$ depends (stochastically) only on the state x and on the action a applied at epoch t (Markov property). As the MDP is assumed to be *stationary*, $p(x'|x, a)$ and $r(x, a)$ are independent of the decision epoch t.

In this paper we focus on a special class of MDP, known as *Graph-based Markov Decision Process* (Chornei et al. (2006)). A GMDP differs from a MDP in that the state and action spaces are multidimensional and that there exist local dependencies between the state and action variables. A GMDP is defined by a five-tuple $< \mathcal{X}, \mathcal{A}, p, r, G >$ where: $\mathcal{X}$ is a Cartesian product of finite sets $\mathcal{X} = \mathcal{X}_1 \times \ldots \times \mathcal{X}_n$; $\mathcal{A}$ is a Cartesian product of finite sets $\mathcal{A} = \mathcal{A}_1 \times \ldots \times \mathcal{A}_n$; p is a transition function $p : \mathcal{X} \times \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$; r is an "immediate" reward function $r : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$; and G is a *directed graph* $G = (V, E)$ expressing *local* dependencies among state and action variables. The local dependencies among the state and action variable make it possible to compute the transitions and rewards from local terms. A *neighbourhood function* N over V is defined as:

Definition 1 (Neighbourhood function) $N : V \rightarrow 2^V$ *is defined as:* $N(i) = \{j \in V | (j, i) \in E\}, \forall i \in V$.

Transitions and rewards are *local* in a GMDP according to the directed graph G:

Definition 2 (Local transitions) *Let* $< \mathcal{X}, \mathcal{A}, p, r, G >$ *be a GMDP. Transitions are* local *iff:*

$$p(x'|x, a) = \prod_{i=1}^{n} p_i(x'_i | x_{N(i)}, a_i), \forall x' \in \mathcal{X}, x \in \mathcal{X}, a \in \mathcal{A}$$

where: $x_I = \{x_j | j \in I\}, \forall I \subseteq \{1, \dots, n\}$.

Definition 3 (Local rewards) *Let* $< \mathcal{X}, \mathcal{A}, p, r, G >$ *be a GMDP. Rewards are* local *iff:*

$$r(x, a) = \sum_{i=1}^{n} r_i(x_{N(i)}, a_i), \forall x \in \mathcal{X}, \forall a \in \mathcal{A}.$$

The action to apply, given the state of the system, is defined through a (stationary) *policy*. A policy is a function $\delta : \mathcal{X} \rightarrow \mathcal{A}$ that assigns at each time epoch an action to every state. In the general case, policies for a GMDP take the form $\delta = (\delta_1, \dots, \delta_n)$, where $\delta_i : \mathcal{X} \rightarrow \mathcal{A}_i$. Nevertheless, global policies can take $O(n\alpha^{\sigma^n})$ space, where $\alpha = \max_i |\mathcal{A}_i|$ and $\sigma = \max_i |\mathcal{X}_i|$. Except for problems of very small dimension, this clearly prohibits the computation of global policies. Some special policies, called *local policies* are therefore of interest:

Definition 4 (Local policy) *Let* $< \mathcal{X}, \mathcal{A}, p, r, G >$ *be a GMDP. A policy* $\delta : \mathcal{X} \rightarrow \mathcal{A}$*, is said to be* local *iff* $\delta = (\delta_1, \dots, \delta_n)$*, where* $\delta_i : \mathcal{X}_{N(i)} \rightarrow \mathcal{A}_i$.

Even though the optimal policy of a GMDP may not be local, it is interesting to look for "good" local policies. This since they are both easy to express (space complexity in $O(n\alpha^\sigma)$) and implement. From now on we will consider only local policies.

Two algorithms for finding approximate solutions to large GMDP have been suggested. One based on Approximate Linear Programming (ALP) (Forsell and Sabbadin (2006)) and one based on Mean-Field Approximation (Peyrard and Sabbadin (2006)). Both methods give an approximate solution to the problem as a local policy, and have linear complexity in the number of variables in the GMDP. The two solution algorithms have been evaluated on small toy problems for which they gave similar results in the value of the policy. We have however in this work only been using the ALP algorithm.

3 GMDP FOREST WIND DAMAGE MODEL

In this section we present the proposed GMDP model. We show how the model is defined and how forest aspects such as stand growth, stochastic wind events, yield from silvicultural management and probability of wind damage are taken into account.

A forest is commonly divided into a set of stands where each stand is a "*geographically contiguous parcel of land considered homogeneous in terms of tree vegetation*" (Lawrence et al. (2001)). In the GMDP model, the state of each stand in the forest is represented by the value of a state variable $\mathcal{X}_i$. If the forest consists of n stand's, the state space of the GMDP is the cross product of the domains of the state variables $\mathcal{X}_i, i = 1, \dots, n$, represents the state of stand i. By representing the state of a stand as the age of the stand, the state variables can both describe the timber value in the stand and the risk of the stand being damaged by wind. Damage to stands due to wind can of course be defined in several different ways. We chose to define a wind damaged stand as a stand that has been so severely damaged by wind that the forest owner will always salvage harvest the stand. More specifically, the state space of each state variable $\mathcal{X}_i$ is the set $\{1, 2, \dots, m, m+1\}, m \in \mathbb{N}$. If the time periods are Y years long, the state of a stand is k and $k < m+1$, then the trees in the stand are between $Y(k-1)$ and $Yk-1$ years old. If the state of the stand is $m+1$, then the stand was damaged by wind between Y and 1 years ago, after which the stand was salvage harvested, and the current trees in the stand are between 0 and $Y-1$ years old. For simplicity, it is assumed that once a stand is sufficiently old (above $Y(m-1)$ years), the state of the stand does not change any more. The timber stock of the stand will no longer grow and its characteristics will not change.

To keep the action space small, only two management activities are free to select, "don't-clear-cut" and "clear-cut". These two management actions are denoted by c_1 and c_2. Besides these two management activities, the stands are also treated according to a number of fixed management activities. These are management activities that will always take place and are not influenced by the two free management activities. The fixed management activities are: site preparation; planting; pre-commercial thinnings; and thinnings. Each stand has an individual specification of the fixed management activities and they are incorporated when the revenues of the stand is calculated.

The state and action spaces of the GMDP model can then be formalised as: I) $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$, where $\mathcal{X}_i = \{1, 2, \dots, m, m+1\}$ and II) $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$, where $\mathcal{A}_i = \{c_1, c_2\}$. The local management policy is a specification of when nothing should be done to a stand or when it should be clear-cut.

For a specific time period, the probability of a stand being damaged by wind depends on the state of the stand and the shelter effect provided by other stands in the forest. This as stands can block the wind and thereby decrease the risk of other stands being damaged. The geographical layout of the forest specifies which stands can provide shelter and which stands they provide shelter to. A directed graph $G = (V, E)$ is used to specify the pattern of sheltering effects. The graph G specifies for each stand i, which stands in the forest influence stand i probability of being damaged by wind. Each stand in the forest is thus represented by a vertex in V. There exists a directed edge $(j, i) \in E$ iff stand j can give shelter to stand i. As the probability of a stand being damaged depends on the state of the stand itself, each node in the graph also has a loop: $(i, i) \in E$.

The time periods are modelled so that damage to the stands due to wind always occurs before any management activity is performed. As it is assumed that a damaged stand will always be salvaged harvested, the selected management activity will never be performed on a stand damaged by wind. That is, given that a stand will be damaged by wind during a time period, the management activities "don't-clear-cut" or "clear-cut" are changed into "salvage-harvest".

The management activities and the damage inflicted by wind affect the normal aging dynamics of the stands and the rewards received. The probability of a stand being damaged by wind is $p_i(x'_i = m+1 | x_{N(i)})$. This transition probability is independent of the action selected, increases with the age of the stand and decreases with the age of the stands that can provide shelter. The aging dynamics of the stands are summarized in Table 1.

Table 1
The aging dynamics of the stands

State	Action	Damaged	Next state
$x_i \neq m+1$	c_1	No	$x'_i = min(m, x_i + 1)$
$x_i = m+1$	c_1	No	$x'_i = 2$
x_i	c_2	No	$x'_i = 1$
x_i	c_1, c_2	Yes	$x'_i = m+1$

As the rewards that are received are different if a stand is damaged or not, the reward functions of the GMDP has to be formulated as $r_i(x_{N(i)}, a_i, x'_i)$. However, the GMDP framework requires them to be locally formulated as $r_i(x_{N(i)}, a_i)$. The previously described reward functions are therefore transformed, leading to an equivalent GMDP where:

$$\hat{r}_i(x_{N(i)}, a_i) = \sum_{x'_i \in \mathcal{X}_i} p_i(x'_i | x_{N(i)}, a_i) r_i(x_{N(i)}, a_i, x'_i),$$
$$\forall x_{N(i)} \in \mathcal{X}_{N(i)}, \forall a_i \in \mathcal{A}_i, i = 1, \dots, n.$$

4 PROBABILITY OF WIND DAMAGE AND REVENUES

The probability of wind damage, as well as revenues generated by the silvicultural operations have to be specified in the GMDP model. For evaluating these, two sub-modules specially developed for the Swedish forestry were used.

4.1 Probability of wind damage

To assess the probability of wind damage, a tool developed and assessed in southern Sweden (Olofsson and Blennow (2005)) was used. As wind damage is usually concentrated at forest edges (Persson (1975)), the tool evaluates a stand's probability of being damaged by wind, through a classification procedure based on the state of the forest and on the geographical location of the stand. More specifically, for a specific edge of a stand, the tool uses a decision tree to classify the edge of the stand as either having a high or low annual probability of wind damage. In this way, all edges of a stand can be classified as either having a high or low annual probability of wind damage. The decision tree we use is taken from (Olofsson and Blennow (2005)) and can be seen in table 2.

The edges that are classified as high probability edges have a yearly probability of being damaged of 5% or more, while low probability edges have a yearly probability of being damaged less than 5%. The decision tree is based on the following characteristics of the forest stand. Stand edge height, the difference between the mean tree height of the stand and its neighbouring stand (m). Stem taper, the ratio of mean tree height to mean diameter at breast height (m/cm). Direction of exposure, the orientation of the edge $\{0°(\text{N}), 60°, 120°, 180°(\text{S}), 240°, 300°\}$. Exp, the topographical shelter of the stand in terms of large and small scale variations of the terrain {exposed, sheltered}. The stem density (no. of stems/ha). For simplicity, we assumed in this study that the landscape was completely flat. The stand edges were therefore never classified as exposed.

Table 2
Decision tree for classification of edges

	Edge height > 10	
	↓	↘
	Stem taper $<= 0.98$	Stem taper > 0.98
↙	↓	↓
Dir. of exposure = 240	Dir. of exposure = 60	**High probability edge**
↓	↓	
Exp = exposed	Stem density $<= 863$	
↓	↓	
High probability edge	Stem taper > 0.82	
	↓	
	High probability edge	

The probability of the stand being damaged by wind was defined as follows. First, the *edge sections* of the stand were specified so that each section was only neighbouring one other stand. Each edge section was then classified as either being a *high impact* or a *low impact* edge section, depending on the length of the edge section and depth of the stand at the edge section. All high impact edge sections were then classified as having a high or low annual probability of wind damage using the above described decision tree. The annual probability of the stand being damaged by wind was then defined as follows. If a least one of the high impact edge sections was classified as a high probability edge, the annual probability of the stand being damaged was defined as 5%. If all the high impact edge sections were classified as low probability edges, the annual probability of the stand being damaged was defined as 0%. This meant that low impact edges had no effect on the probability of a stand being damaged.

4.2 Revenues

The revenues generated by performing the silvicultural management actions in the stands were specified through production tables. These tables were generated by a growth-and-yield simulator (Wikström (2000)), and one table was generated for each stand. Each table specifies the revenue generated by clear-cutting or salvage harvesting the stand. It also specifies the revenues generated by the fixed and underlying management regimes.

Two additional forest aspects had to be taken into account in the specification of the silvicultural revenues: I) Thinning is never performed in the same time period as a clear-cut and II) according to the Swedish Forestry Act it is illegal to clear-cut young stands. The first aspect was taken into account by constructing the revenue from clear-cutting in a time period so that it does not include any revenue from the fixed management activities. The second constraint was handled by setting the revenue from clear-cutting to a large negative number in case a stand was clear-cut too early. By setting this cost sufficiently high, no stand would be clear-cut too early. For a summary of the revenues, see Table 3.

Table 3
Revenue generated by a specific action and damage to the stand

Action	Damaged	Revenue
Don't-clear-cut	No	Fixed management
Clear-cut	No	Clear-cut
Don't-clear-cut	Yes	Salvage harvest
Clear-cut	Yes	Salvage harvest

5 PRELIMINARY RESULTS

To evaluate the suggested method, it was tested on a forest estate in southern Sweden called Björnstorp (55°37′N / 13°24′E) which is located within the Swedish temperate zone. The estate amounts to a total of 2800 ha of which around 1200 ha is forest mainly dominated by Norway spruce. This is the same estate as the one used in (Olofsson and Blennow (2005)) to build the decision tree described in 4.1. As a preliminary study only a section of the Björnstorp estate was considered. The selected area contained 39 forest stands to a total of 99 ha. We selected time periods of 20 years. The GMDP model then contained 39 forest areas (n=39), each stand could take six values (ages), and two actions could be performed in the stands ($|X_i| = 6, |A_i| = 2, i = 1, \ldots, 39$). An annual discount factor of 2% was used and all results were generated on an Intel Pentium 3.6 GHZ/ 1.00 GB machine using a Scilab implementation.

In the selected section of the estate, only fourteen stands had a stochastic development. For the other stands there were never any risk of wind damage as there risk levels were lower than the thresholds used in this study. The edge sections of these stands were always classified as having a low annual probability of wind damage. To assess the value of taking the risk of wind damage into account, two silvicultural management policies were compared. A "storm" policy taking the risk of wind damage into account and a "no storm" policy not taking the risk of wind damage into account. The storm policy is the management policy found by solving the GMDP model of the estate with the ALP algorithm. The no-storm policy is the policy found by applying the ALP algorithm to a no risk version of the GMDP model of the estate in which all risks of wind damage had been removed. That is, the probability of a stand being damaged by wind was always equal to zero. The two management policies were evaluated on the original GMDP model of the estate with standard Monte Carlo simulations of 400 trajectories of length 20 time periods and where the current state of the forest was used as the starting point. The expected NPV (10^3 SEK) of the estate when it is managed according to the storm and no-storm policies were 1323 and 1303. The expected NPV of the areas having a stochastic development when they are managed according to the storm and no-storm policies were 124 and 104.

6 DISCUSSION

We have described a work in progress concerning the management of a forest faced with risk of wind damage. We have suggested a GMDP model that can be used to optimise the silvicultural management policy when the risk of wind damage is taken into account. We have shown how the model is specified and how it takes into account stochastic wind and growth events as well as several important aspects of wind damage. The probability of wind damage is modelled as depending on the state of the stands, the state of the neighbouring stands, the spatial structure of the forest, the geographical orientation of the forest as well as the topography of the area. The suggested model is also suitable for modelling large forest areas as the complexity of the ALP solution algorithm only grows linearly in the number of stands.

Preliminary tests on a case study show that the suggested model can be used to find a silvicultural management policy that takes the risk of wind damage into account. Even though only a third (14/39) of the stands in the case study were facing risks of wind damage according to the criteria used, the value of the forest increased by managing it according to a policy taking the risk of wind damage into account. The expected NPV of the whole forest only increased slightly by 2%. However, for the stands at risk of wind damage, the expected NPV increased by as much as 19%. This shows that for stands at risk of being damaged by wind it is important to find a management policy that takes into account the risk of wind damage.

In this model we have assumed relatively low risk levels of the stands. We have modelled the minimal risk levels according to the tool used for characterising the risk levels of the edges of the stands. It is likely that the risk levels of the stands are substantially higher than these. As the estimation of the expected NPV of the forest is based on the risk levels, it is likely that an increase in the risk levels will further increase the value of managing the forest according to the storm policy. Such an analysis is however left as a future work but the proposed model is a useful tool for such an analysis.

Considering the preliminary results on the Björnstorp estate, further developments of the model will be made. When the neighbouring graph in the GMDP model was created only geographical constraints were taken into account. Two stands were defined as neighbours if they share a border. However, some edge sections are always classified as having a low annual probability of wind damage. If both edge sections are always classified as having a low annual probability of wind damage the two stands have no influence on each other's development and should not be linked in the GMDP graph. As a result, some stands are completely independent of their neighbours and their optimal management policy can easily be computed as it only depends on the state of the stand. The size of the GMDP model may in this manner be dramatically reduced, in the considered section of the Björnstorp estate 25 of the 39 stands can be computed individually. The same remark applies to the computation of the no-storm policy that was computed for evaluation of the suggested method. The policy was computed by solving a GMDP with independent vertices. When the risk of wind damage is not considered the stands are independent and individual and optimal management policies can easily be computed. This in comparison to the ALP solution method that in theory is approximate.

Further developments of the suggested model will also be concentrated on the development of solution algorithms for GMDP based on Reinforcement Learning (Sutton (1991)). As a GMDP can be seen as a multi-agent collaborative MDP, these solution algorithms will be in line with the ones suggested by (Guestrin et al. (2002), Kok and Vlassis (2006)). We hope that such solution algorithms will make it possible to shorten the time periods used in the model. Shorter time periods would improve the management options of the stands as well as the possible interactions between the stands. This would increase the management possibilities for the stands at risk. Such improvements would give us further knowledge of the value of taking the risk of wind damage into account in forest management policies.

7 REFERENCES

Chornei, R. K., H. Daduna, and P. S. Knopov (2006). *Control of Spatially Structured Random Processes and Random Fields with Applications*. Springer.

Forsell, N. and R. Sabbadin (2006). Approximate linear-programming algorithms for graph-base markov decision processes. In *Proc. 17th European Conf. on Art. Int. (ECAI'2006)*, Riva Del Garda, Italy, pp. 590–594.

Gardiner, B. A., G. R. Stacey, R. E. Belcher, and C. J. Wood (1997). Field and wind tunnel assessments of the implications of respacing and thinning for tree stability. *Forestry 70*, 233–252.

Guestrin, C., M. G. Lagoudakis, and R. Parr (2002). Coordinated reinforcement learning. In *ICML*, pp. 227–234.

Heinonen, T. and T. Pukkala (2004). A comparison of one- and two-compartment neighbourhoods in heuristic search with spatial forest management goals. *Silva Fennica 38*, 319–332.

Kellomäki, S. and H. Väisänen (1997). Modelling the dynamics of the forest ecosystem for climate change studies in the boreal conditions. *Ecological Modelling 97*, 121–140.

Kok, J. R. and N. A. Vlassis (2006). Collaborative multiagent reinforcement learning by payoff propagation. *Journal of Machine Learning Research 7*, 1789–1828.

Lawrence, S. D., K. N. Johnson, P. S. Bettinger, and T. E.Howard (2001). *Forest Management : to sustain ecological, economic, and social values, 4th Edition*. Mc Graw Hill.

Lohmander, P. and F. Helles (1987). Windthrow probability as a function of stand characteristics and shelter. *Schandinavian Journal of Forest Research 2*, 227–238.

Meilby, H., N. Strange, and B. J. Thorsen (2001). Optimal spatial harvest planning under risk of windthrow. *Forest Ecology and Management 149*, 15–31.

Olofsson, E. and E. Blennow (2005). Decision support for identifying spruce forest stand edges with high probability of wind damage. *Forest Ecology and Management 207*, 87–98.

Peltola, H., S. Kellomäki, H. Väisänen, and V. Ikonen (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*.

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*.

Peyrard, N. and R. Sabbadin (2006). Mean field approximation of the policy iteration algorithm for graph-base markov decision processes. In *Proc. 17th European Conf. on Art. Int. (ECAI'2006)*, Riva Del Garda, Italy.

Pukkala, T. and M. Kurttila (2005). Examining the performance of six heuristic optimisation techniques in different forest planning problems. *Silva Fennica 39*, 67–80.

Puterman, M. L. (1994). *Markov Decision Processes*. New York: John Wiley and Sons.

Quine, C., M. Coutts, B. Gardiner, and G. Pyatt (1995). *Forest and wind: Management to minimise damage*. London: Bulletin 114, HMSO.

Schelhaas, M., G. Nabuurs, and A. Schuck (2003). Natural disturbances in the european forests in the 19th and 20th centuries. *Global Change Biology 9*, 1620–1633.

Schroeder, L. and H. Eidmann (1993). Attacks of bark- and wood-boring coleoptera on snow-broken conifers over a two-year period. *Schandinavian Journal of Forest Research 8*, 257–265.

Sondell, J. (2006). Erferenheter från "operation gudrun". *Skogforsk Resultat nr 7*.

Sutton, R. (1991). Planning by incremental dynamic programming. In M. Kaufmann (Ed.), *Proc. of the 8th International Workshop on Machine Learning*, pp. 353–357.

Valinger, E. and J. Fridman (1997). Modelling probability of snow and wind damage in scots pine stands using tree characteristics. *Forest Ecology and Management 97*, 215–222.

Valinger, E. and N. Pettersson (1996). Wind and snow damage in a thinning and fertilization experiment in picea abies in southern sweden. *Forestry 69*, 25–33.

Wikström, P. (2000). A solution method for uneven-aged management applied to norway spruce. *Forest Science 46*, 452–462.

Zeng, H., T. Pukkala, and H. Peltola (2007). The use of heuristic optimization in risk management of wind damage in forest planning. *Forest Ecology and Management 241*, 189–199.