WORKING PAPER
1998-8

sustainable
forest
management
network

réseau
sur la
gestion durable
des forêts

A Grid-Based Spatial Model of Forest Dynamics Applied to the Boreal Mixedwood Region

A Network of Centres of Excellence
Un réseau de centres d'excellence

S. Cumming, D. Demarchi and C. Walters

For copies of this or other SFM publications contact:

Sustainable Forest Management Network
G208 Biological Sciences Building
University of Alberta
Edmonton, Alberta, T6G 2E9
Ph: (780) 492 6659
Fax: (780) 492 8160
http://www.ualberta.ca/sfm

A grid-based spatial model of forest dynamics applied to the boreal mixedwood region.

S. Cumming[1], D. Demarchi[2] and C. Walters[2]

[1]Department of Renewable Resources, University of Alberta
[2]Department of Zoology, UBC

August 31, 1998

1 Introduction

This report describes a landscape simulation model being developed for applications to forest management and habitat conservation problems in the boreal mixedwood forest. The overall modelling framework is described, with particular emphasis on the steps being taken to tailor the model for use in the mixedwood region of Alberta. The regionally-specific model elements are fire, mixedwood stand dynamics, hierarchical harvest scheduling, and habitat modelling. These components are described in varying detail, depending on the state of the underlying research programmes. Some evolving applications of the model are briefly noted. This work is in progress. All design decisions, model components and process parameters should be expected to undergo change in the future.

We are starting from an existing computer simulation model that was developed to provide an objective and quantitative tool for comparing forest harvesting and conservation policies for the spotted owl and other rare species (Demarchi, 1998). The model consists of an individual-based population dynamics model which is linked to a detailed spatial simulation of forest harvesting, to represent as accurately as possible the forest landscape changes that a particular species of wildlife will face. The model has been used to evaluate alternate forest management policies aimed at protecting and conserving the northern spotted owl on timber supply areas in the vicinity of Vancouver, B. C. The model software is written in Visual Basic 5.0 under Microsoft Windows 95.

The model is based on a simulation modelling environment developed by C. Walters over the past decade. The environment is designed for the rapid development of spatial dynamic models of ecological processes which can be used either as research tools, or in an interactive "workshop" environment. The user-interface in particular has been developed to facilitate the rapid generation and visualisation of alternate scenarios. Versions of the model have been applied to a wide range of management, policy and ecological problems, in both terrestrial and marine ecosystems.

Early versions of Walter's model ran at low spatial resolutions. In some applications, harvest scheduling was limited to clearcut logging at a 2500 ha resolution. This was clearly undesirable, and as technology developed, the potential for higher resolution modelling was realised. In the spotted-owl application, a spatial resolution of 25 ha and spatial extent of $\approx 2.3 \times 10^6$ ha. This resolution approaches the extent and grain of actual forestry practices. The mixedwood version discussed here is being developed at a 9 ha resolution, and will be able to simulate areas as large as several million ha. The rationale for this is given in Section 2, below.

2 Data Layers

The model reads and operates on raster-based map files, such as can be generated from GIS data. The area of concern and the desired spatial resolution must be defined beforehand, as part of the data assembly phase. Once entered, data layers can be manipulated through a user-friendly interface, either before or during a simulation run. This feature is especially useful in modelling workshops. In forestry applications, data layers typically include

1. location of the forested cell;
2. elevation;
3. forest type and age;
4. site productivity indices;
5. road access and potential haul-routes;
6. reserve status (*i.e.*, parks and conservation areas) and
7. known animal locations (for population modelling applications).

2.1 Study area and resolution

We used ARC/INFO coverages of AVI data from part of the Alberta-Pacific Forest Management Area (FMA) as the primary data source. We selected a 3×3 township region ($\approx 8.4 \times 10^4$ ha) from FMU S8 as our pilot study area. After considerable experimentation, where we compared the ability to detect features in the gridded landscapes against the storage costs of higher resolution, we chose 9 ha as our grid size. Although larger than the minimum mapping unit in the underlying inventories (1 ha), 9 ha:

1. is close to the mean size of all common stand types in Phase 3 inventories (Cumming *et al.*, 1996);
2. is smaller than most operationally feasible cut-block sizes;
3. approximates the area-of-influence of small patches of white-spruce (Stewart *et al.*, 1998); and
4. will permit us to exploit most of the available data regarding forest fragmentation effects in the mixedwood Schmiegelow *et al.* (1997).

Although fires smaller than 9 ha can not be simulated, these apparently account for less than 0.001 of total area burnt by lightning fires in northeastern Alberta. We expect to be able to model areas up to several million ha at this resolution.

2.2 Forest type and age

Forest type and "age" are obtained from digital forest inventory maps, which are gridded down to the chosen spatial resolution and then classified into a small number of pre-determined cover-type classes. We classed AVI polygons using a simple scheme that experience has shown to be useful: 1) water; 2) muskeg; 3) aspen; 4) mixture of aspen and white spruce; 5) white spruce; 6) black spruce; and 7) pine. Age of forested types was inherited from the AVI stand-origin field. A revised classification scheme for the mixedwood component was developed later, as part of the stand dynamics submodel (Section 5, below). This new scheme requires several subsidiary data layers, which describe stand trajectories and regeneration substrate.

2.3 Other data layers

In previous applications, elevation has been used to constrain the dynamic road building submodel, as costs become very high in mountainous terrain. Elevation and other topographic parameters may be used as part of habitat suitability models or to limit *e.g.,* dispersal by individual animals. They may also have some utility in the fire and stand dynamics components. Our group has developed a variable resolution digital elevation model for northeastern Alberta, which could be gridded down for use in this model, and at least one FMA-holder (Daishowai-Marubeni International) has elevation, slope and aspect attributes tied to each AVI polygon. However, at present, we do not plan to use topographic data in our simulations, since it is not yet a factor in any of our underlying statistical models of forest dynamics. Similarly, spatial climate data, though available, have not yet been included. These data layers can easily be incorporated when required.

Existing provisions for incorporating site productivity indices are a holdover from B. C. applications. In AVI, site classes are defined as ranges of site index values. They could therefore be be used to constrain yield curves and developmental trajectories of regenerating stands, in later versions of the model.

At this stage in development, we are not concerned with the actual location of road networks. This information can be added at any time.

3 Harvest Scheduling

The model generates timber harvest schedules based on a number of economic and operational criteria. Resultant changes in forest structure can then be used to assess relative responses of other model processes over time via the component submodels of fire, stand dynamics, wildlife habitat and population dynamics. The most important aspect of the forest harvesting component is that it acts at a resolution and precision comparable to that of the underlying digital forest inventory data (9 ha). Cut-levels and operability, scheduling and blocking rules (how cells are aggregated into larger harvestable units) can be combined to create a rich array of harvesting strategies and local tactics.

3.1 Dynamic road building

Haul costs are an important consideration in any logging plan. By finding optimal log hauling routes from each forest cell to the nearest mill, the model can locate and cut the timber based not only on its characteristics, but on whether or not it is economical to harvest. The procedure is relatively simple in concept. Starting out with known mill locations, the model recursively finds, for each cell, the cheapest way to build road access to and haul wood from that cell to a mill. This process continues until all the forest cells in the database are directly linked to a mill. At each stage, for each unprocessed cell, the minimum cost hauling directions for all surrounding cells are checked, and the new cell is linked to the neighbour with the lowest cost-to-mill, to minimise road building and transport costs.

Setting connections without regard to geographical contours or mill location could result in hauling routes that are illogical (*e.g.,* over top of high-elevation passes, long and/or adverse hauling routes). To reflect this, the cost of establishing a new link is related to elevation changes. Optimal haul routes follow elevational contours that minimise road-building costs. The cost parameters for these decisions are all specified by the user. In mixedwood applications, it may be that cover type (*e.g.,* muskeg) should be a more important consideration than topography. The cost parameters may be specified and altered by the user.

Once the dynamic program has assigned a hauling direction to each forest cell, the linkage structure is saved, This eliminates the need to initialise the network program at the start of each simulation, which reduces startup time. It is necessary to update the road linkage as access develops during harvesting, and previously unroaded cells become “roaded.”

At present, roads access is computed to grid-cell resolution which is unnecessarily high for the present application. If the dynamic road-building feature proves useful in the mixedwood context, later versions will reduce the granularity of the road consruction and access to something more reasonable (*e.g.,* townships).

3.2 Cost accounting

The cost and production values (in dollars) that result from harvesting are estimated from a few key parameters. From these, and the given access state, harvesting and haul costs are determined. Once a block is chosen for logging, these costs plus cutting costs and net production value (price per m^3 × logged volume) are calculated. The cost of harvesting a cell is then the cost per km of building a road to that cell from the nearest road, plus the cost of cutting the timber, plus basic silviculture cost, plus the cost of hauling wood back to the nearest mill location (cost/km x distance). Net production value is the total price of annual wood production less total annual costs. During a simulation, the user may plot the total area logged, costs and production per time. This allows visual interpretation of patterns such as "fall-down" and changes in operating costs as less accessible areas are harvested.

3.3 Harvest block layout

Our 9 ha grid cells are smaller than the average cut-block sizes used in Alberta. Therefore, adjacent cells are aggregated into larger, relatively homogeneous areas called *blocks*. Blocks may be either operable or inoperable. Inoperable blocks are those that are in reserves, are too old or too young to cut, contain never-merchantable forest types, or that are under some other restriction. Operable blocks are the model's harvestable units.

The present version of the model automatically blocks cells, attempting to match a given target size (*e.g.,* 100 ha) and within-block age variability. Both these parameters can be set by the user (see the menu items in the upper left corner of Figure 2). An example block layout is shown in Figure 1. In the model version now being developed, block sizes follow a log-normal distribution of user-specified mean and variance: this family of distributions yield good fits to the actual size distributions observed in inventory data (Cumming *et al.*, 1996). In addition, the within-block variability in conifer content will be specifiable.

3.4 Cut sequence

The basic parameter for cut sequencing is the Annual Allowable Cut (AAC) which provides a volume target for the scheduler. AAC is a regional total, not a per-destination value. In the case of multiple delivery sites, allocation of harvest to specific locations (mills or export points) is controlled entirely by the haul-cost calculations. The basic unit of harvest is the block. Six block selection options are supported:

1. sequential
2. oldest blocks first;
3. highest net volume;
4. the maximum net value (*i.e.* production minus cost);
5. random
6. two-pass harvest within townships or similar areal units.

Figure 1: Sample of automatic cut-block layout, with block size 60 ha and within-block age range of 30 yr. Not all blocks are distinguishable in this grey-scale image.

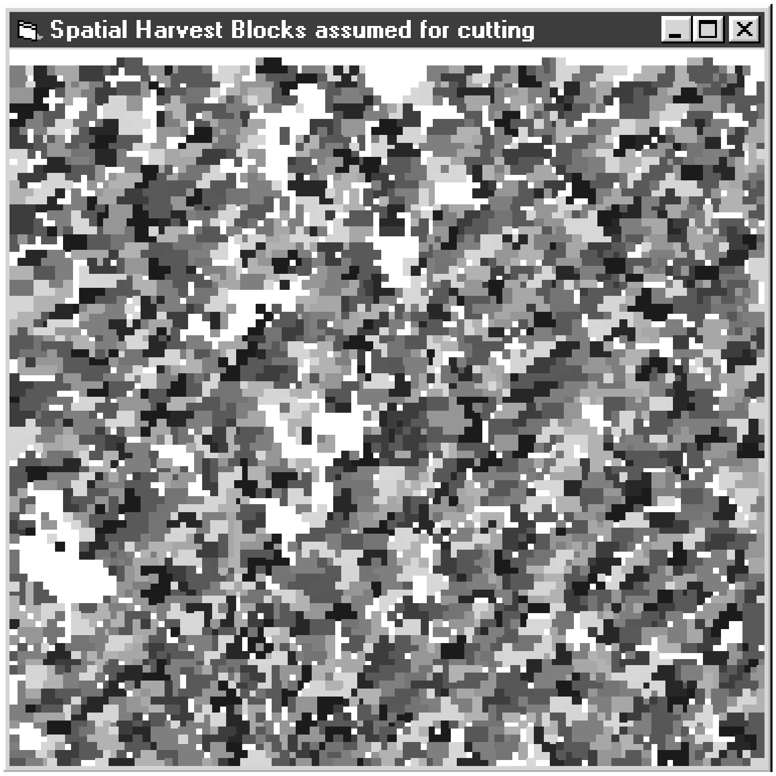

For cases 2–5, all operable blocks are ranked by the chosen criterion, and top ranking cells are harvested until the AAC is reached. Random mode means what it says. For the *township/2 pass* strategy, the area is first gridded down into townships (these, or other harvestable units, may also be specified as an input data layer). Then, each block within the township is scheduled such that at most half of the blocks may be cut in a single time step. The parameter `Yr to 2nd pass` controls the interval between consecutive township entries. In the version now under development will be more fully hierarchical, in that within-township block selection will follow one of the three ordering criteria listed above. The results of two sample schedules following widely different strategies (two-pass and sequential) are shown in Figure 3.

Since blocks can contain cells of widely different ages, some cells may not be operable when the block is selected, and these are not cut (see the `Forest Harvest Parameters ...` portion of Figure 2). Extended green-up periods and/or strategies to minimise visual impacts, are simulated by establishing temporary reserves in adjacent cells within a user-specified radius of the cut. These reserves are marked by a

Figure 2: Sample harvest parameter form

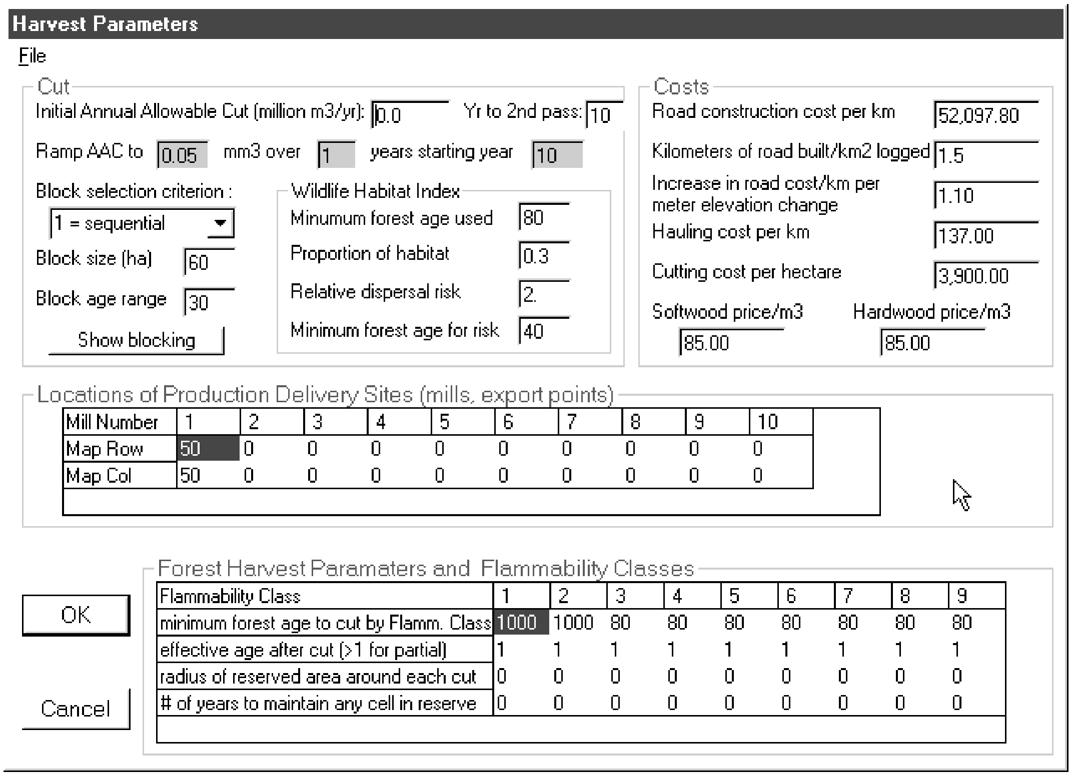

Cut

Initial Annual Allowable Cut (million m3/yr): 0.0

Yr to 2nd pass: 10

Ramp AAC to 0.05 mm3 over 1 years starting year 10

Block selection criterion : 1 = sequential

Block size (ha) 60

Block age range 30

Wildlife Habitat Index

Minumum forest age used 80

Proportion of habitat 0.3

Relative dispersal risk 2.

Minimum forest age for risk 40

Costs

Road construction cost per km 52,097.80

Kilometers of road built/km2 logged 1.5

Increase in road cost/km per meter elevation change 1.10

Hauling cost per km 137.00

Cutting cost per hectare 3,900.00

Softwood price/m3 85.00

Hardwood price/m3 85.00

Locations of Production Delivery Sites (mills, export points)

Mill Number	1	2	3	4	5	6	7	8	9	10
Map Row	50	0	0	0	0	0	0	0	0	0
Map Col	50	0	0	0	0	0	0	0	0	0

Forest Harvest Paramaters and Flammability Classes

Flammability Class	1	2	3	4	5	6	7	8	9
minimum forest age to cut by Flamm. Class	1000	1000	80	80	80	80	80	80	80
effective age after cut (>1 for partial)	1	1	1	1	1	1	1	1	1
radius of reserved area around each cut	0	0	0	0	0	0	0	0	0
# of years to maintain any cell in reserve	0	0	0	0	0	0	0	0	0

flag which is decremented at each time step. When the flag reaches zero, the reserve or greenup condition expires[1].

Simulations may be paused at at any point, to allow the user to change the values of particular parameters. By displaying a `Forest Harvesting Parameter Form` (Figure 2), the user can dynamically manipulate such factors as the block selection rule, AAC level, green-up requirements and economic parameters.

The model tracks the specific cells within each block that are cut in each time step, as well as various aggregate statistics (*e.g.*, total area cut). These may be saved to a `Harvest Scenario File` to provide simulated disturbance histories for population models (Section 6.4).

---

[1] Permanent reserves are dealt with simply by making their expiration date long enough. Permanent or dynamic reserve networks are beyond the scope of this report.

Figure 3: Two extremal harvest scheduling strategies: two-pass, within township (top) and progressive clearcut, (bottom). Both show the state of the forest after 100 years of harvesting at 50, 000 $\mathrm{m}^3\ \mathrm{yr}^{-1}$.

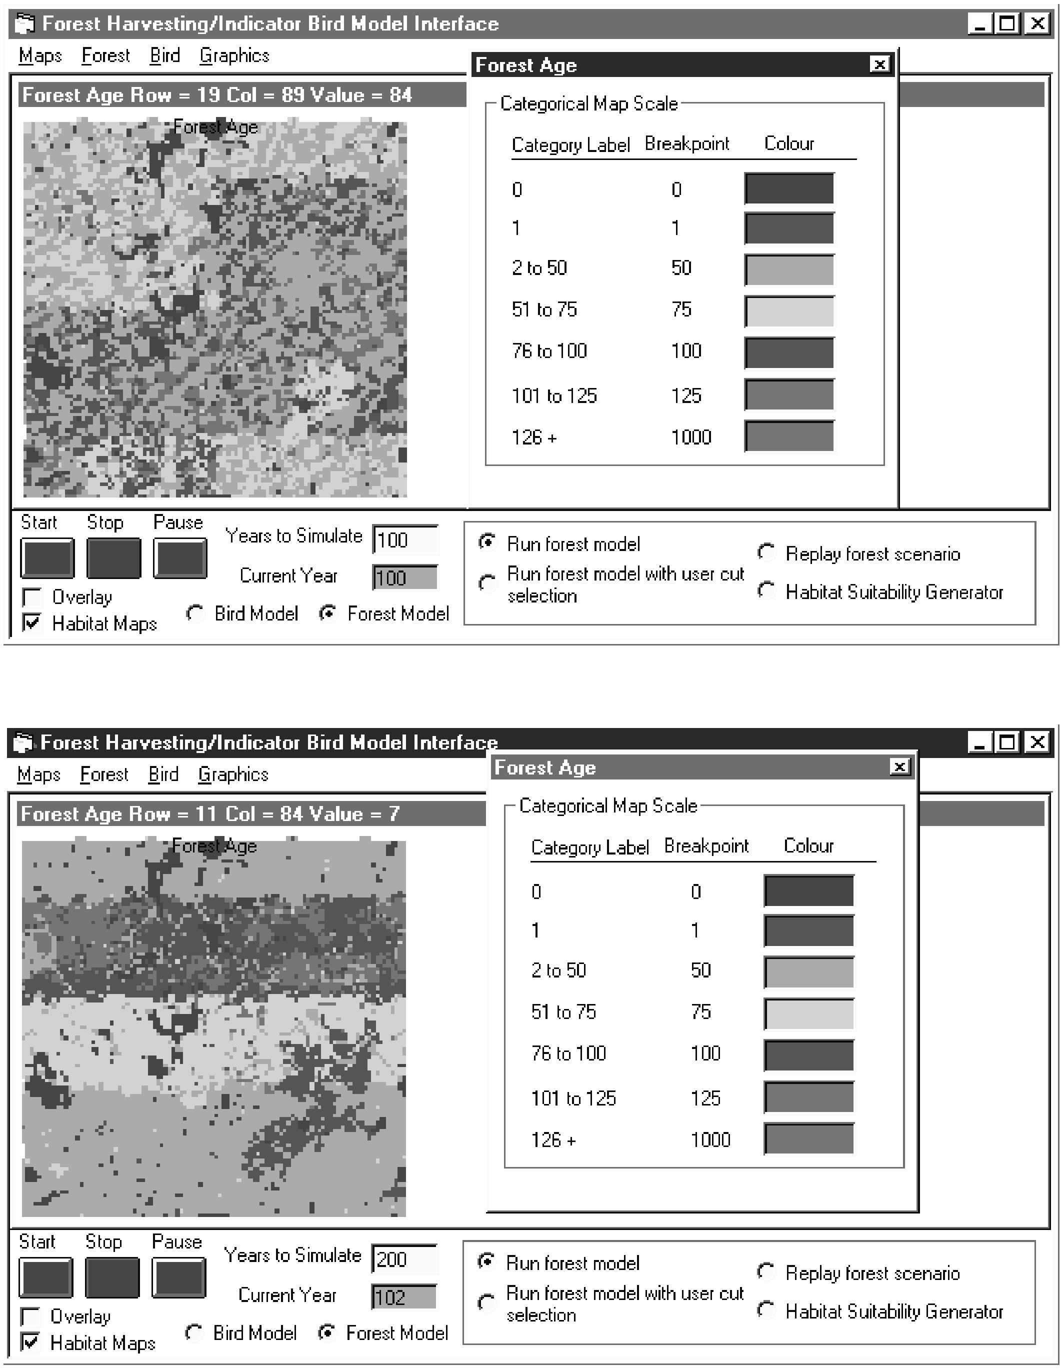

Forest Age Row = 19 Col = 89 Value = 84

Forest Age

Categorical Map Scale

Category Label	Breakpoint	Colour
0	0	■
1	1	■
2 to 50	50	■
51 to 75	75	□
76 to 100	100	■
101 to 125	125	■
126 +	1000	■

Years to Simulate 100

Current Year 100

☐ Overlay
☑ Habitat Maps
○ Bird Model ● Forest Model

● Run forest model
○ Run forest model with user cut selection
○ Replay forest scenario
○ Habitat Suitability Generator

Forest Age Row = 11 Col = 84 Value = 7

Forest Age

Categorical Map Scale

Category Label	Breakpoint	Colour
0	0	■
1	1	■
2 to 50	50	■
51 to 75	75	□
76 to 100	100	■
101 to 125	125	■
126 +	1000	■

Years to Simulate 200

Current Year 102

☐ Overlay
☑ Habitat Maps
○ Bird Model ● Forest Model

● Run forest model
○ Run forest model with user cut selection
○ Replay forest scenario
○ Habitat Suitability Generator

3.5 Yield data

A set of almost 200 yield tables and corresponding inventory was provided by Alberta-Pacific Forest Industries Inc., in support of a related modelling initiative. Most combinations of species composition, stand density and site had a specific table. This precision is unwieldy for the purposes at hand. An aggregated yield table was assembled for each of the five forested types described in Section 2.2. This was done by classifying the initial set of tables down to those basic types, and then taking an area-weighted mean of yield table entries, based on the abundances of each initial type in the received inventory. A similar translation will be performed for the new mixedwood model of Section 5, but the details are not yet determined.

3.6 Silviculture

At present, only a limited set of silvicultural treatments are simulated. Logging is followed by natural regeneration. The determination of initial trajectories for regenerating stands is discussed in Section 5. More complex submodels representing *e.g.,* partial cuts and planting may be added later, as interest from our industrial partners warrants, and as available data and resources permit.

4 Fire

In this version of the model, we consider lightning caused fires only. The model treats fire as a three stage stochastic process. Stage 1 models the ignition process, and determines if a fire actually starts in a particular grid cell in a given model time step. Stage 2 models the effect of fire suppression, in effect by determining a probability that a fire stays within the cell of origin. Stage 3 models fire growth, using spread parameters or "jump probabilities" which determine the probability of a fire spreading across a cell boundary, from a burning cell into an unburnt neighbour. These three processes depend on parameters that can be set by the user, using the menu screen shown in Figure 4.

Fires are started by applying a random number sequence to each cell independently. Fires are extinguished when no further spread events (jumps) occur. After a fire, forested cells are reset to age 1. Thus, fires are presumed to be stand initiating. These processes result in irregular patches of variously sized burns. The locations of fires are not constant across simulations because of stochastic patterns of fire ignition and growth. As with harvesting, the model tracks the sizes and compositions of each burn. Fire sizes follow a probability distribution largely determined by the spread parameter(s). In this section, we briefly describe the steps taken so far to estimate parameters for these three processes.

4.1 Ignition probabilities

To a first approximation, fire ignition may be treated as a spatial Poisson process, where the number of fires (n) starting in a given area per unit time has the probability density function

$$p(X = n) = \frac{\lambda^n}{n!} e^{-\lambda}$$

The parameter λ can be estimated by counting the number of observed fires in a given area over a given time interval: that is, the parameter is an index of the density or frequency of ignitions per unit area, per unit time. For the Alberta-Pacific FMA, over the interval 1968–1993,

$$\lambda_a = 3.332$$

in units of $10,000 \cdot 26 \text{ ha}^{-1} \text{ yr}^{-1}$. That is, the mean number of fires expected over 26 years, per $10,000$ ha is 3.332. This value was fairly easy to estimate from datasets available from Alberta Environmental Protection (AEP). The lower time limit of 1968 was chosen to reduce detection bias, since the network of fire towers was complete by that year. The upper limit of 1993 reflects the year (1994) when a consistent data base was assembled by Cumming.

Figure 4: Sample fire parameter screen, with our current parameter estimates.

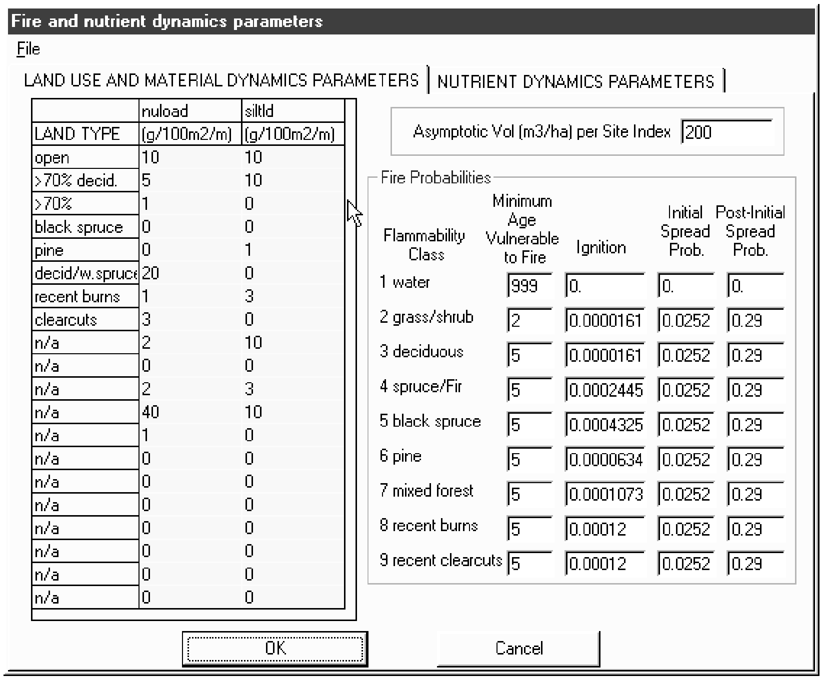

	nuload	siltld
LAND TYPE	(g/100m2/m)	(g/100m2/m)
open	10	10
>70% decid.	5	10
>70%	1	0
black spruce	0	0
pine	0	1
decid/w.spruc	20	0
recent burns	1	3
clearcuts	3	0
n/a	2	10
n/a	0	0
n/a	2	3
n/a	40	10
n/a	1	0
n/a	0	0
n/a	0	0
n/a	0	0
n/a	0	0
n/a	0	0
n/a	0	0
n/a	0	0

Asymptotic Vol (m3/ha) per Site Index 200

Fire Probabilities

Flammability Class	Minimum Age Vulnerable to Fire	Ignition	Initial Spread Prob.	Post-Initial Spread Prob.
1 water	999	0.	0.	0.
2 grass/shrub	2	0.0000161	0.0252	0.29
3 deciduous	5	0.0000161	0.0252	0.29
4 spruce/Fir	5	0.0002445	0.0252	0.29
5 black spruce	5	0.0004325	0.0252	0.29
6 pine	5	0.0000634	0.0252	0.29
7 mixed forest	5	0.0001073	0.0252	0.29
8 recent burns	5	0.00012	0.0252	0.29
9 recent clearcuts	5	0.00012	0.0252	0.29

To scale the previous result to a 9 ha grid cell at an annual resolution is easy, assuming a Poisson process:

$$\lambda = \lambda_a \frac{9}{26 \cdot 10,000} = 0.000115$$

The probability of more than a single fire starting per cell in a given year is given by

$$p(X > 1) = 1 - p(X = 0) - p(X = 1) \approx 10^{-8}.$$

So, λ may be treated as the annual ignition probability per cell, to a very good approximation, and the possibility of multiple ignitions can be ignored at this spatial and temporal scale.

4.1.1 Fuel type effects: ignition

In northeastern Alberta, the assumption of pure spatial randomness in ignition is not valid. At the township scale ($\approx$10,000 ha), the frequency of ignition is strongly related to township *composition* as measured by the proportion of the township area covered by following 6 exhaustive types: leading pine, leading black spruce, aspen dominated (D and D(C) stands, white spruce dominated (C and C(D) stands) mixed

Figure 5: Relative frequencies of fueltype at initial attack, compared to expectation based on the compositions of the recorded townships of ignition.

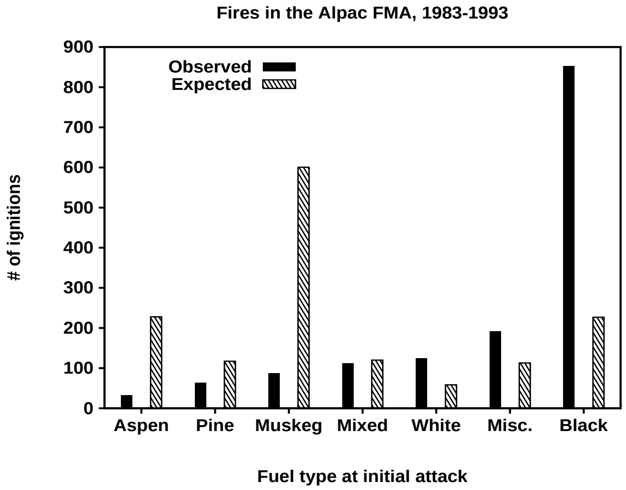

Type	Count	Expected	Factor	λ_t
Aspen	33	228.0	0.145	1.66e-5
Pine	64	117.2	0.546	6.28e-5
Muskeg	87	600.6	0.145	1.67e-5
Mixed	112	120.0	0.934	1.07e-4
White Sp.	124	58.4	2.122	2.44e-4
Misc.	192	113.0	1.699	1.95e-4
Black Sp.	852	226.9	3.755	4.32e-4

Table 1: Actual counts *vs.* expected number of initial fueltype assessments, their ratio (Factor) and estimated annual ignition probability for each type (λ_t).

stands (CD and DC stands) and "other" which is mostly various kinds of muskeg. These are the same categories used in the present model. One can explain more than 50% of the deviance in ignition density using these variables: but so far, only at the scale of townships and $\approx$25 yr intervals (S. Cumming,unpublished data.[2]). Attempts to scale these results down to 9 ha grid cells and an annual time step have so far failed. Depending upon the precise statistical model used, either some less-flammable

---

[2] See `http://www.rr.ualberta.ca/research/BMDP/ignite.pdf`

types are predicted to have non-physical, negative ignition probabilities or else some rare but flammable types (notably, pure white spruce stands) have implausibly high ignition probabilities. The reason for these difficulties appears to be lack of contrast in the underlying datasets: there are no townships with high proportions of white spruce. Because spatial variability in ignition is liable to be a very important factor in explaining the structure of the forest, and in modelling the interaction between harvesting and fire, a less satisfactory approach has been adopted for the time being.

Fire fighting crews assess the *fueltype at initial attack* for each fire (that is fought). This attribute is recorded in the Provincial fire databases. From 1983-1993, the attribute can be translated directly to Phase 3 forest inventory classes. Our present results use an approximate translation, although the "official" one has been provided by the staff of Forest Protection, AEP, which will be used for the next iteration of model development. If ignition is random at the scale of an individual lightning strike then: the probability of a fire starting in a given fueltype should be the proportional area of that fueltype in the area of the fire. We use the composition of the township in which the fire was detected to estimate this proportion. Figure 5 shows the expected *vs.* observed frequencies of the "initial fueltype" attribute, for 1,464 fires in 548 townships from 1983–1993. Clearly, some types are more frequently encountered than would be expected. Assuming that "fueltype at initial attack" is equivalent to "fueltype at ignition," Table 1 presents another view of the incidence data, and shows how fueltype-specific ignition parameters have been derived from the average value of $\lambda = 0.000115$. Even for the most fire-prone type (Black spruce), the probability of multiple ignitions per time step is still small ($\approx 9 \times 10^{-8}$), so the λ_t may be treated as ignition probabilities.

4.2 Fire Suppression

Fire suppression seems mostly to act by keeping fires smaller than about 2 ha (Cumming, 1997). This suggests a way to model the initial stage of fire growth, incorporating the effect of varying degrees of fire suppression effort. We divide the fire growth model into two components, "initial spread" and later growth. Let p_0 be the probability that a real detected fire remains smaller than a single grid cell (< 9 ha). This can be estimated empirically from fire history data. In the model, for a fire to have final size < 9 ha, $n = 8$ independent spread events (one for each neighbouring cell) must fail to occur. If the probability of initial spread across an edge is p_i, then the probability of *not* spreading is $(1 - p_i)$ and we have

$$p_0 = (1 - p_i)^n,$$

or

$$p_i = 1 - \exp(\log(p_0)/n).$$

Interval	Ignitions	p_s	p_i
1983–1993	1467	0.953647	0.0059
1971–1982	1211	0.877787	0.0162
1961–1970	260	0.819231	0.0246

Table 2: Empirical initial spread probabilities for three time intervals, derived from Alberta Environmental Protection databases for the Alberta-Pacific FMA.

Table 2 shows the raw counts, probabilities and computed p_i values for three fire-suppression regimes which can be understood as corresponding to full, intermediate and minimal levels of suppression effort (Cumming *et al.*, 1995).

4.2.1 Fuel type effects: initial spread

Perhaps the initial spread probabilities should be refined so as to be cell-type (that is to say, forest-type) specific. There may be enough replicates of different initial types to permit a logistic regression analysis. However, differences in the size of the fire at initial attack will significantly complicate the analysis. We will examine this problem later.

4.3 Spread probabilities

In estimating jump probabilities (p_j, the fire spread parameter) the key thing is to match the mean size and the frequency of occurrence of large fires in the empirical data. In this case, we used the same fire history data set mentioned in Section 4.1. The first-approximation estimate for the p_j was obtained by comparing the simulated and empirical size distributions for fires (> 9 ha). We generated sets of sample fires on the study area landscape, using a range of jump probabilities from 0.20 to 0.40 (Table 3). Based on various statistical tests, it appears that the correct value is approximately 0.30. Although $p_j = 0.35$ produces the best overall fit to the empirical size distribution on visual inspection (Figure 6) and according to a K-S test, this value produces too many large fires, which inflates the mean size. The Anderson-Darling test, which is sensitive to differences between to distributions at their extrema, detects the overabundance of large fires.

The estimated p_j was further refined by "replaying" historical fires. These were interpreted as larger patches in the inventory data composed of 40–60 yr old stands. By dividing the number of actual jumps taken by the number of tests made, the value of p_j necessary to produce the observed fires can be estimated. This method leads to an estimated $p_j = 0.29$. An instance of a large fire, simulated under the parameters described in this section, is shown in Figure 7. Such large fires seem

Table 3: Summary statistics of fire size distributions simulated using various values of the jump parameter p_j. Sizes are in units of 9 ha pixels. "n" is sample size, $\bar{x}$ is sample mean, $x_{\max}$ is sample maximum and x_{95} is the approximate 95th level of a variance-corrected Student's t-test, p_{KS} is the significance level of a Kolgomorov-Smirnov and AD is the value of the Anderson-Darling test-statistic obtained by comparing the simulated and empirical size distributions (row "True")

p_j	n	$\bar{x}$	$x_{\max}$	x_{95}	p_t	p_{KS}	AD
0.20	4031	17	185	57	0.006	$< .0001$	1424
0.25	3351	44	629	183	0.022	$< .0001$	423
0.30	1763	157	18737	641	0.77	$< .0001$	123
0.35	551	691	28030	295	0.005	0.85	473
0.40	199	2125	30351	29672	< 0.001	0.0029	6304
True	263	175	11521	492	na	na	na

to occur about once or twice a century on the model landscape. The simulation generates fire islands, which suggests a further way to calibrate model behaviour, using the results of Eberhart and Woodard (1987).

4.3.1 Fuel type effects: fire spread

In reality, fire spread across a cell boundary may depend on the nature of the boundary being crossed. For example, a fire may be more likely to spread into a black spruce stand than into an aspen stand. That means that jump probabilities should be type-specific, It is also possible that the probabilities are non-symmetric (*i.e.* edge-type specific.) Suitable empirical data to directly estimate these probabilities is scarce.

We are developing an approach to this problem based on some statistical models that relate the composition of a fire to the composition of the locality where it started (Cumming, 1997). We envision a fairly complex combination of simulation and maximum likelihood methods based on these models. The details are beyond the scope of this report.

Figure 6: Cumulative size distributions for fires generated with jump probability p_j = 0.25, 0.30 and 0.35, compared with empirical data from the Alberta-Pacific FMA.

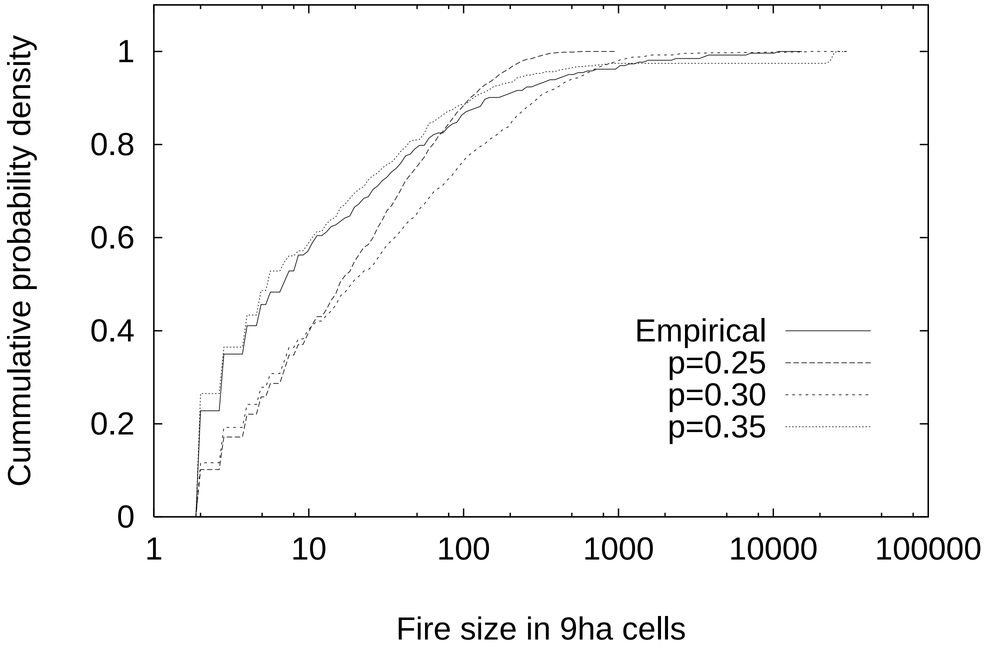

Figure 7: A large fire (Category Label 1, bottom right of map) which burnt within a predominantly old forest. Water bodies are light grey (Category 0). Other forest types are mapped to the same grey-scale value, to emphasise the fire.

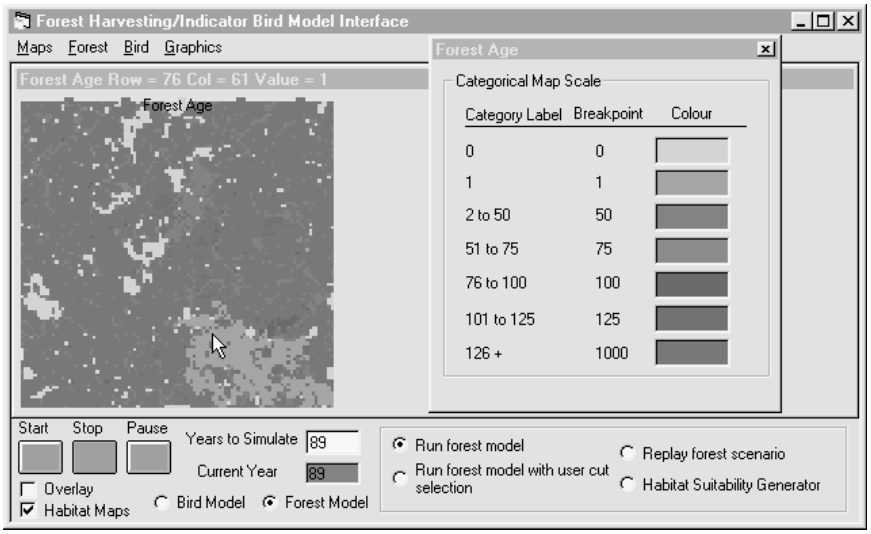

Forest Age Row = 76 Col = 61 Value = 1

Categorical Map Scale

Category Label	Breakpoint	Colour
0	0	▨
1	1	▨
2 to 50	50	▨
51 to 75	75	▨
76 to 100	100	▨
101 to 125	125	▨
126 +	1000	▨

Years to Simulate 89
Current Year 89
☐ Overlay
☑ Habitat Maps
○ Bird Model
● Forest Model
● Run forest model
○ Run forest model with user cut selection
○ Replay forest scenario
○ Habitat Suitability Generator

5 Mixedwood Dynamics

This section describes a model of the stand dynamics of the mixedwood component of the landscape, designed *de novo* for incorporation into our landscape simulator. This component is crucial for correct representation of fire ignition and spread, studying interactions between harvesting and fire, and projecting the future distribution and availability of white spruce.

Our objective was to define a biologically reasonable model of the dynamics of mesic stands, using as few parameters as possible, and which implicitly makes use of allometric relationships such as that between volume and seed production. We have attempted to keep the model dimensionless. Most components (seed production, recruitment, etc.) are described by variables or functions which take values between 0 and 1. We believe most of the empirical claims and assumptions can be supported by the current literature on white spruce ecology and aspen stand dynamics.

Mesic stands are considered to be those that can support both *Populus* species (primarily trembling aspen) and white spruce: this status is inferred from inventory data. The principal state variable describing a model cell is N_t, which measures the proportion of total standing volume at time t that is composed of white spruce. We may assume that the balance $(1 - N_t)$ of the stand volume is aspen. As a stand develops after initiation, the relative dominance of spruce can increase, from an initial value which is almost always close to zero. We are concerned to model proportional spruce *volume* because the stratification of mesic stands used to develop our statistical models of fire ignition is essentially volume-based. When absolute volumes of merchantable timber are required for the purposes of the present model, they can be estimated from yield tables and the other state variable, t, which measures stand age. t in the current model is time since disturbance by fire or clearcut logging.

The basic idea is that a mesic stands follows a trajectory (N_t, t) defined by a logistic growth curve

$$N_t = f(t; N_0, r, K)$$

defined by the three parameters N_0, the initial spruce content, r, the intrinsic rate at which spuce content increases, and K, the final proportion of spruce towards which the stand is tending (see Figure 8). For example, a trajectory with $K = 0$ describes a stable stand of pure aspen. The initial spruce composition of a mesic stand will almost always be very low, because of early establishment and dominance by aspen, which is almost always present. The initial trajectory of the stand depends on random factors, such as fire and seed production. The model is dynamic in that a stands trajectory can change over time as a result of recruitment through seed input from adjacent stands, or of local origin as the initial cohort of white spruce (if any) matures. This is modelled by changing the parameters of the logistic growth model, subject to some constraints, as described below.

N_t will suffice to classify mesic stands into fire ignition and fire spread classes.

Figure 8: Dynamics of N_t (proportional volume of white spruce) on a well-stocked "mixedwood" stand (AwSw) on a fair quality site, as given by simulated yield tables, and a "fit-by-eye" logistic growth curve.

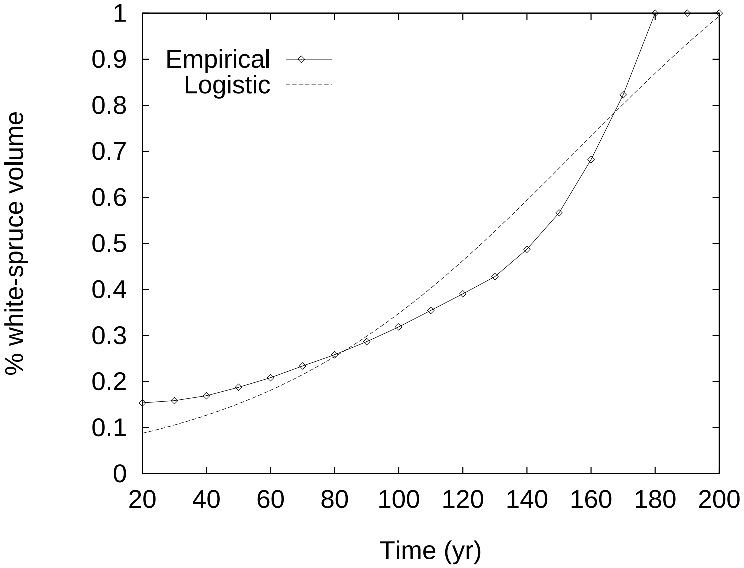

This model will replace the current static stratification of the forest into "pure aspen," mixedwood and "pure spruce." Other stand types we are considering (*e.g.,* pine, black spruce and wetlands) will not be effected, and can continue to be represented as they are. This version of the model represents as few phenomena as possible. In particular it does not include mortality: the assumption is that stands follow a trajectory which eventually stabilises, and remains there until either external conditions change or the stand burns.

5.1 Seed production

The intuition here is that a mature, fully stocked white spruce stand is capable of producing enough seeds to fill all available substrate in the vicinity, at least in a masting year, and some lesser proportion q_s in other years. Assume that mast years occur with probability $m \approx 0.2$. We assume that mast years occur synchronously at the spatial extent which we are modelling ($\approx 100,000$ ha),

Define I_t to be local seed production at time t. I_t has units of potential regeneration *and survival to the canopy* per unit of available suitable substrate. The seed production model is

$$I_t = f(N_t, t)$$

for mast years, and

$$I_t = q_s \cdot f(N_t, t)$$

otherwise. f can be chosen as some allometric function relating seed production to spruce density N_t and absolute time t. However, the function $f(N_t, t) = N_t$ will work just fine, for now at least. Thus, a fully stocked mature stand of spruce ($N_t = K_t = 1.0$) has seed production $I_t = 1.0$ in mast years. This sub-model requires 2 global parameters m and q_s which can be estimated from the literature. It requires no additional state information.

5.2 Seed dispersal

The model at present uses 9 ha cells, which represent areas 300 m a side, or from centre to centre. This is approximately the 90% dispersal range for white spruce seeds (Stewart *et al.*, 1998). Let us assume that a cell's production I_t is allocated spatially as follows:

1. a proportion p_l is allocated locally, within the cell.
2. a proportion $q_l = 1 - p_l$ is equally divided amongst the cells 8 distance-1 neighbours.

Define the total seed inputs to a cell at time t as

$$I_t^* = p_l I_t + \sum_j q_l I_{t,j}$$

where the summation is over the 8 distance-1 neighbours. Stewart *et al.* (1998) contains results from which to estimate p_l. For the present, assume $p_l = 0.5$. Then, maximum seed production in all eight neighbours produces a total exogenous seed input of 0.5. It may also be worth examining a longer-tailed dispersal function, going out to distance-2 or distance-3 neighbours, with appropriate rarefaction. It is almost obvious that $I_t^* \leq 1$, regardless of distance considered and the decay constant p_l, but to be safe, this should be enforced in the model code.

In its distance-1 form, this submodel requires essentially one parameter (p_l) which can be estimated from the literature, and no additional state information. It can be implemented by a simple filter.

5.3 Substrate model

White spruce recruitment is substrate limited. Suitable substrate is mainly either exposed mineral soil, or decayed downed logs sites. Mineral soil is exposed by intense fire, but its availability declines rapidly as the site is colonised. Downed log input

Figure 9: Dynamics of white-spruce seed bed availability, subsequent to an intense fire at $t = 0$.

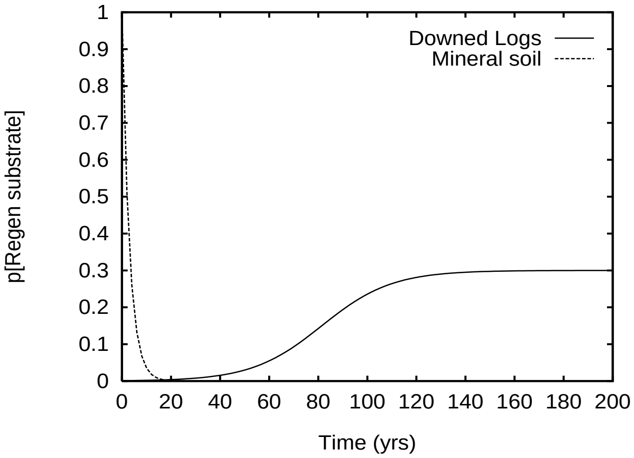

requires the development of a suitable overstory and subsequent mortality, which we assume to equilibrate at some value as the stand ages.

Substrate dynamics is modelled by

$$S_t = m_t + d_t$$

where m_t is mineral soil availability and d_t is downed log substrate. As with the seed production model, these two terms are scaled so that $0 \leq S_t \leq 1$. Figure 9 illustrates the dynamics: process of initial rapid decline from full availability, as exposed mineral soil is rapidly colonised by plants, and the gradual development of autogenically derived substrate as the stand matures. The dynamics of mineral soil availability (m_t) post fire is modelled by

$$m_t = s_0 * e^{-t/b}$$

where $s_0 \leq 1.0$ and the decay parameter $b = 3.0$. Variation in fire intensity could be incorporated by making s_0 a random variable, with support $[0, 1]$ Some studies by E. MacDonald at the University of Alberta and/or by K. Miyanishi[3] may lead to an underlying distribution from which to sample. In the meantime, assuming $s_0 = 1 - d_0$ will probably do, where d_0 is the initial density of downed-log substrate. We may be able to estimate b from the published work of Johnson and by Lieffers, who have

---

[3] Reported at an NCE Regeneration Umbrella workshop, held at Montreal, April 1998.

inferred changes post-fire white spruce recruitment rates from stand age-structures. The value of 3 shown is probably a bit low (V. Lieffers, personal communication).

The later development of substrate (d_t) within the cell depends on autogenic disturbance processes, mainly the death and decay of canopy trees. We model this as a sigmoid function of stand age. In Figure 9, post-fire substrate development follows

$$d_t = \frac{d_0 k e^{rt}}{k + d_0(e^{rt} - 1)}$$

with $d_0 = 0.001$, $r = 0.07$ and $k = 0.3$. Analysis of the Calling Lake vegetation data (F. Schmiegelow) and the results of Lee *et al.* (1997) should yield estimates of downed log input rates and decay rates, respectively, while Lieffers *et al.* (1996) will yield estimates of recruitment rates. Thus, extant data will provide a fire approximation of parameter estimates for this submodel.

This component requires a total of 5 parameters, and since it depends only on time, requires no additional state variables. The implicit assumption is that downed log input and subsequent decay do not depend on the composition of the stand, *i.e.* on N_t. When N_t approaches 1.0, so that most input would be spruce logs, then species-specific mortality and decay rates may begin to be significant. But, by then, the stand is full of spruce, and little recruitment can occur.

5.4 Recruitment

Recruitment into a cell modifies the cell's trajectory by increasing the end-point density of spruce, or the "carrying capacity." The dynamics are:

$$R_t = (1 - K_t)^{\alpha} \cdot I_t^* \cdot S_t$$

where K_t is the current end-point, I_t^* is the seed input and S_t is the substrate availability, and $\alpha \geq 1$ is a competition coefficient which models the effects of canopy composition on spruce recruitment. As a default, $\alpha = 1$. However, K. Statd has reanalysed some data from Lieffers *et al.* (1996), and determined that height growth of understory spruce is negatively related to the amount of spruce in the canopy. Further work will suggest a value for α and/or a revised functional form for the competition term.

The condition

$$0 \leq R_t \leq (1 - K_t)$$

holds, since the other terms take values between 0 and 1. This means that recruitment can, at most, cause the eventual full occupation of the site. If K_t is close to one, even a combination of high seed input and good substrate availability will have little effect. However, in a small patch just after a fire ($S_t \approx 1$), high seed input from neighbouring mature spruce cells could cause the establishment of nearly pure spruce stand, at least in a mast year. This component requires no additional state variables, and at most the additional parameter α.

5.5 Stand Dynamics

Stand trajectories

$$N_t = f(t; N_0, r_t, K_t)$$

have the same logistic form as the downed log substrate model,

$$N_t = \frac{N_0 K_t e^{r_t t}}{K_t + N_0(e^{r_t t} - 1)}$$

except that K_t and r_t depend on time. This is the crucial point, because this time dependency is how changes in the stand trajectory are modelled.

The initial stand trajectory is determined by post-fire recruitment R_0, which establishes the stand on a trajectory partly determined by

$$K_0 = R_0$$

r_0 will be set so that growth rate is maximised at about the time white spruce volume increment is maximised, or about $t = 80$ —this can be determined from the literature or from yield tables we have from various sources. We have not yet determined how to set N_0.

The stand trajectory can change through later recruitment, which acts by increasing the maximum white spruce dominance value:

$$K_{t+1} = K_t + R_t$$

In reality, the effect of recruitment on standing crop and growth rates is lagged. By increasing K, it is necessary to also change some other parameters of the logistic growth curve. In order that recruitment not change the instantaneous growth rate of the stand, r_{t+1} is chosen so that df/dt is constant. That is,

$$r_t N_t(1 - N_t/K_t) = r_{t+1} N_t(1 - N_t/K_{t+1})$$

Write $K_{t+1} = K_t + dk$ and $r_{t+1} = r_t + dr$. Then $dr < 0$ is given

$$dr = \frac{[N_t(1 - N_t/K_t)r_t] - [N_t(1 - N_t/(K_t + dk))r_t]}{N_t(1 - N_t/(K_t + dk))}$$

This first order condition on r is intended to approximate the behaviour of a more complex age-structured model. In order that recruitment not alter standing biomass, we also redefine N_0 according to:

$$N_0 = \frac{N_t * K_t}{K_t \exp(r_{t+1} t) - N_t(\exp(r_{t+1} t) - 1)}$$

Figure 10: Changes in spruce trajectory with a recruitment pulse at year 50, under various re-parameterisations, which are explained in the text.

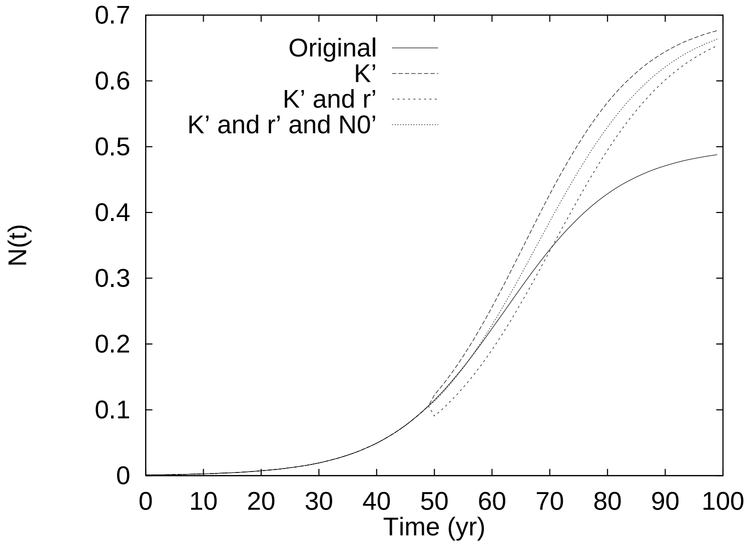

This post-hoc change in the previous trajectory is admittedly a bit arbitrary. Figure 10 illustrates the effect of these changes in parameters. We begin with $N_0 = 0.001$, $r_0 = 0.1$ and $K_0 = 0.5$, and add a pulse of recruits $R_{50} = 0.2$ at year 50. The updated parameters are $r_{50} = 0.0928$ and $N_0 = 0.00129$. The reader may care to verify these results against the stated formulae.

A number of studies being conducted by members of our group will probably yield more defensible parameter estimates and validation data (*e.g.*, the distribution of white spruce dominance in stands of known age). It may also be that some stand-level growth and yield experts can suggest an alternate functional form for the spruce-dominance model. Also, note that the present model description is playing fast and loose with notation, and using the continuous form of the logistic to describe what is really a discrete logistic model.

5.6 Model integration and validation

Linking this component into the landscape model still requires that we:

1. Represent silviculture;
2. Link stand trajectories to yield data;
3. Initialise the model from inventory data.

At the time of writing, we have not yet decided how to do this.

This submodel is essentially a population dynamics model for white spruce. The reverse calibration method outlined in Section 6.4 is therefore applicable, especially since the present location of at least mature white spruce stands is pretty well known.

This model represents most of the factors thought to control the distribution of white spruce in the boreal mixedwood, with two exceptions: site factors and individual seed trees. Site factors are partly accounted for by the prior classification into mesic and non-mesic stand types, based on AVI typing. However, it is probable that spruce recruitment is limited by site factors on some sites presently dominated by aspen. We are also aware of the potential importance of large individual white spruce trees as seed sources—some supporting data are reported by Stewart *et al.* (1998). The number of such trees necessary to produce sufficient seed to restock an area, or at least saturate available substrate, is probably below the detection threshold of current inventories. Neither of these two effects can be simulated using available data.

6 Habitat and population models

This section describes our current habitat and population modelling capabilities, and briefly outlines how the model will be linked to the results of the empirical habitat models being developed by Schmiegelow, Vernier and other members of our group.

6.1 Individual-based models

A major limitation of past habitat models is their coarse spatial scales and the resulting inclusion of habitat not normally used by the species of concern. In reality, it is nearly impossible to calculate accurate indices of habitat quality at scales coarser than that of typical forest stands (*i.e.*, between 10–100 ha. Existing models that represent animal territories or discrete habitats do so as symmetrical shapes (*e.g.,* squares, hexagons). Animals do not choose their habitat in this way, and pretending that they do can lead to problems. First, subdividing the landscape up into small symmetric shapes increases the risk of including habitat that is not normally used by, or is inaccessible to, the animals (*e.g.,* non-forested land, territories that cover movement barriers). Second, symmetrical territories allow more animals to fit into the landscape than is feasible, resulting in an overestimate of carrying capacity. This overestimate results from the systematic placement of territories on the landscape. That is, the total area of the population's range divided by the size of an average territory equals the maximum number of animals that could exist if each territory was occupied. This does not allow for variation in the shape or size of habitats due to geographic features and constraints. Another weakness common to habitat models is their inability to realistically represent changes to the habitat base (*e.g.,*

Figure 11: Sample parameter screen for individual-based models.

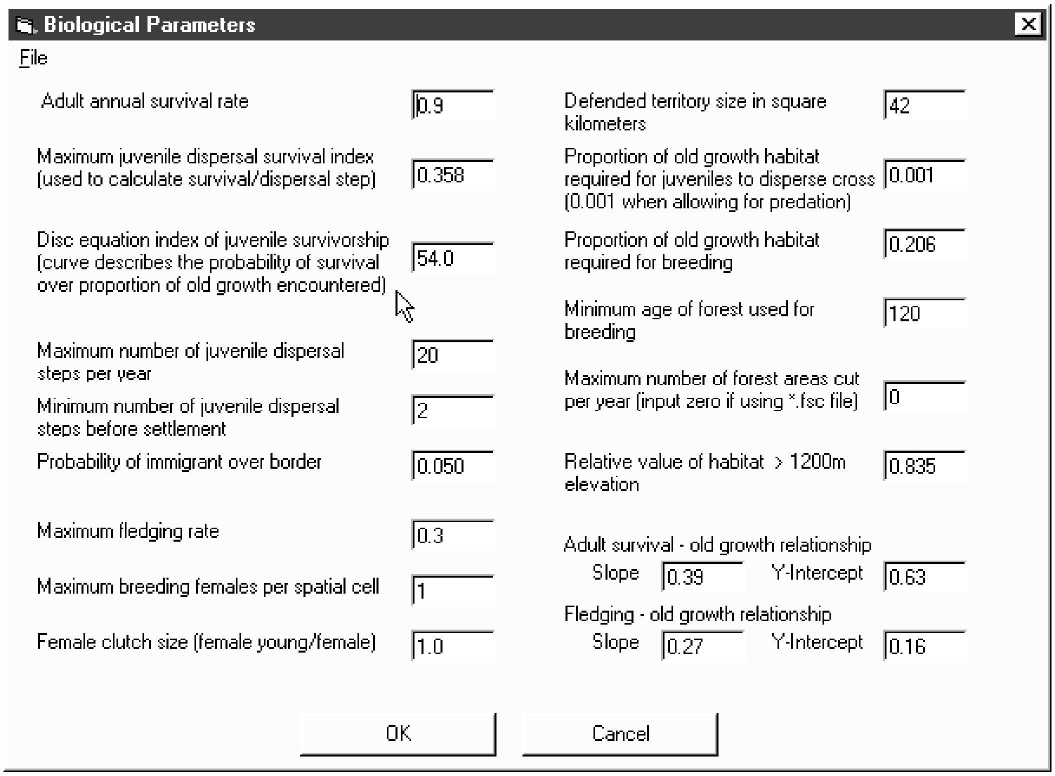

Adult annual survival rate	0.9
Maximum juvenile dispersal survival index (used to calculate survival/dispersal step)	0.358
Disc equation index of juvenile survivorship (curve describes the probability of survival over proportion of old growth encountered)	54.0
Maximum number of juvenile dispersal steps per year	20
Minimum number of juvenile dispersal steps before settlement	2
Probability of immigrant over border	0.050
Maximum fledging rate	0.3
Maximum breeding females per spatial cell	1
Female clutch size (female young/female)	1.0
Defended territory size in square kilometers	42
Proportion of old growth habitat required for juveniles to disperse cross (0.001 when allowing for predation)	0.001
Proportion of old growth habitat required for breeding	0.206
Minimum age of forest used for breeding	120
Maximum number of forest areas cut per year (input zero if using *.fsc file)	0
Relative value of habitat > 1200m elevation	0.835

Adult survival - old growth relationship
Slope 0.39 Y-Intercept 0.63

Fledging - old growth relationship
Slope 0.27 Y-Intercept 0.16

habitat change is either random or abrupt with new habitat maps being inserted at predetermined intervals).

Population simulation methods typically incorporate difference equations based on population size. These assume that the single state variable representing population size can be determined by aggregating individual members of the population. Thus the models are expected to average over the full distribution of individual types in the population (DeAngelis and Gross, 1992). Population modelers recognise that individuals are unique, but think that they be described by statistical parameters estimated from the entire population. DeAngelis and Gross (1992) also pointed out that models without spatial-temporal dependence assume that the environment and every member of the population affects every other member of the population equally. These assumptions violate the principals of locality and of individual variability. One response to this perceived inadequacy has been the development of age- or size-class structured models. A second approach, that avoids aggregation altogether, describes populations by simulating large numbers of individuals simultaneously. For example, a model can imitate the movements and fates of individual animals based on the habitat that they encounter. Both of these alternative approaches have been referred to as individual-based modelling but have been distinguished as individual state (i-state) distribution and i-state configuration models, respectively Metz and Diekmann (1986).

The use of individual-based models to make or defend management decisions is becoming increasingly common in conservation biology. They offer the possibility of better understanding of population behaviour and underlying mechanisms (Botsford, 1992). To realise this potential the models must be realistic and their structure must be comparable to the observed structure of the population itself.

Our objective is to define some useful indices or measures of biodiversity response to the interaction of alternate harvesting strategies with natural disturbance and stand dynamics. At present, the model is set up to look at impacts on an "owl-like" indicator species with large territories and possible dispersal bottlenecks. The model simulates the movements and fates of individual animals according the local habitat configuration in light of the best available data. Changes in habitat over time and the spatial location of each individual determine its fate. The model is necessarily stochastic so results vary from simulation to simulation. Thus, we use Monte Carlo trials to compare outcomes under different scenarios, and to allow visualisation of variability and uncertainty about population trajectories. The model tracks, displays and logs the total number of animals each year for each simulation trial. It also keeps track of means and standard deviations of population levels.

Individuals are characterised by home range requirements and life history parameters such as fecundity, and survival rates. Populations persist in virtue of dispersal of recruits, as they search for territories. Potential territories are assembled analogous to harvest block formation. Mortality risks are associated with dispersal, related to distance travelled and the nature of habitat traversed. All these parameters can be modified by the user (see, *e.g.,* the menu items labelled `Wildlife Habitat Index` in Figure 2) and the full life-history screen of Figure 11.

6.2 Cross-Scale Linkage

We have sought new ways to deal with the spatial cross-scale problem of how to simultaneously represent forest dynamics on a fine spatial scale while representing owl population dynamics on a coarser territory “grid”. The key to this cross-linkage methodology has been to develop a method for mapping potential animal home ranges over the detailed forest landscape data, thus providing a coarse-scale mosaic of animal-use areas. The model then aggregates forest-state data (*i.e.* the suitability measure) over these animal-use areas to provide indices of habitat quality. Each territory is assigned a quality index by adding up habitat index values over the cells contained in the territory. For the territory to meet the life requisite requirements, then its index must meet or exceed the required user-defined index. For old-growth dependent species, the required index is calculated by multiplying the proportion of old-growth required for breeding by the number of cells per territory and then adding the number of cells per territory to that product. For example if a territory is comprised of 25 cells (25 x 9 ha = 225 ha.) and a minimum of 20% old-growth within the territory is required then the minimum suitability index would be 30 (*i.e.*, 25 cells x 0.20 + 25 cells = 30). This calculation effectively omits the contribution of young forested cells to the index. It ensures the index is based solely on the contribution of old forest and that animals only occupy territories with a suitability index ≥ 30. The territory suitability index is updated each year during the course of a simulation. The logging/conservation scenario that was passed from the forest harvesting component to the population component is used to change the forest ages over time. Territory scale suitability indices are updated each year to reflect these changes.

6.3 Habitat suitability models

The model computes edge length indices of total length of edges between old and young stands (this might be a good predictor of habitat quality for some edge-associated species like deer). There is potential to develop a suite of indices that better reflect some species life history or behavioural characteristics (*e.g.,* forages only in old areas, uses edges but subject to higher predation risk there, exploits large territories to average over small scale variation). The model also has the capability to calculate a variety of landscape metrics at variable patch sizes (*e.g.,* length of forest edge, proportion of mature forest, proportion of young forest, proportion of 6-30 year old clearcuts, road density, proportion of non-forest land). These metrics are used by user defined regression models that predict species abundance or presence/absence or some other suitability measure—a prototype menu screen for defining these models is shown in Figure 12. The result is a habitat suitability map that is updated yearly as the landscape simulation progresses. This allows visualisation and quantification of habitat changes. Other members of our group are working up the Calling Lake datasets assembled by Schmiegelow and others, to develop empirically based habitat models for forest songbirds, which will be linked with this modelling framework.

6.4 Historical Forest Disturbance

The model can be used to identify parameter combinations that would allow the simulated population to terminate at current population estimates, given a history of forest disturbance in the study area. A reconstruction of the historical forest disturbance can be used to simulate population dynamics over the last 100 years. Using the historical changes in habitat configuration, the model can then be used to determine whether the model's predicted outcomes (*e.g.,* population level trends, current population size) match the actual response of the population. Such *reconstructive*

Figure 12: Prototype menu for specifying species-specific logistic regression models.

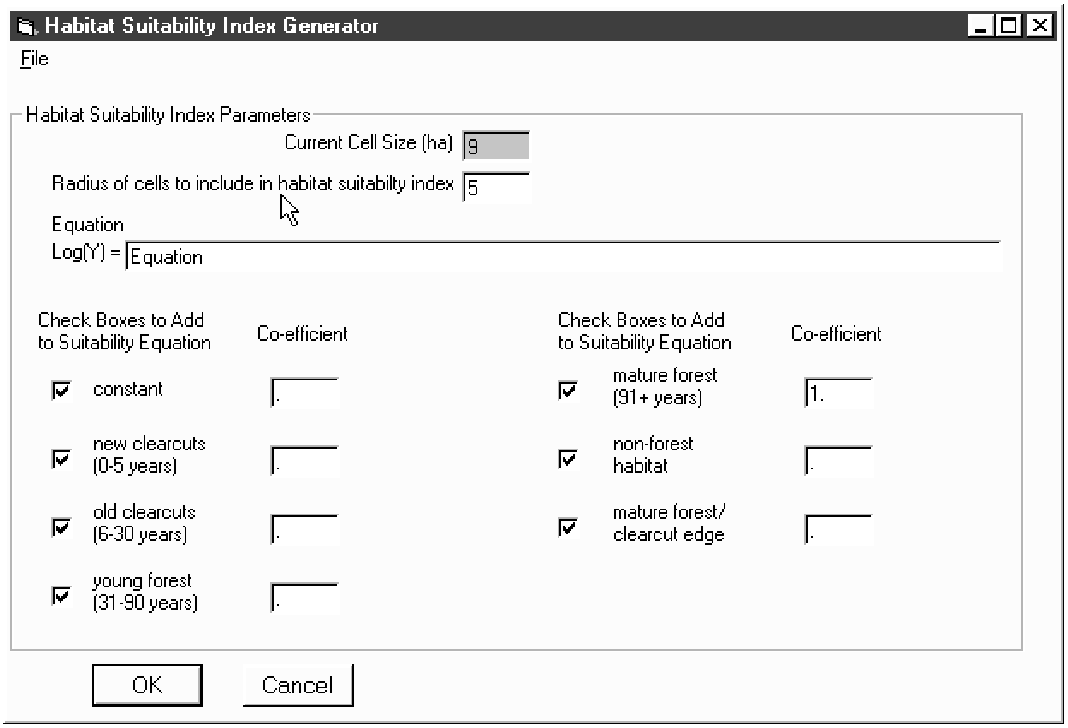

Current Cell Size (ha) 9

Radius of cells to include in habitat suitabilty index 5

Equation
Log(Y) = Equation

Check Boxes to Add to Suitability Equation	Co-efficient	Check Boxes to Add to Suitability Equation	Co-efficient
☑ constant	.	☑ mature forest (91+ years)	1.
☑ new clearcuts (0-5 years)	.	☑ non-forest habitat	.
☑ old clearcuts (6-30 years)	.	☑ mature forest/ clearcut edge	.
☑ young forest (31-90 years)	.		

simulations can be a powerful check on model parameter values. Where records are unavailable, the historical disturbance regime may be inferred from the current forest age data by taking the current age of each cell and calculating when it was disturbed by development, burning, logging or agriculture. In the BC Spotted Owl application, such information was compiled for each stand disturbed within the last 100 years and converted to a forest harvest scenario file which was used in the individual-based owl model. Bart (1995) suggests that this type of "reality check" may perhaps the best evaluation of a model's predictions.

6.5 Applications

Demarchi is working with several University of Alberta researchers to apply our individual based modelling capabilities to two boreal species: the barred owl, and the three-toed woodpecker. The latter case involves some interesting potential interactions between wildfire and dispersal. Potential applications to woodland caribou populations are being explored by Demarchi, Walters and Rohner, a researcher for the West-Central Caribou Committee.

References

Bart, J. (1995). Acceptance criteria for using individual-based models to make management decisions. *Ecological Applications*, **5**, 411–420.

Botsford, L. W. (1992). Individual state structure in population models. In D. L. DeAngelis and L. J. Gross, editors, *Individual-based models and approaches in ecology*. Chapman and Hall, New York, New York.

Cumming, S. G. (1997). *Landscape Dynamics of the Boreal Mixedwood Forest*. Ph.D. thesis, University of British Columbia, Vancouver, B. C., Canada.

Cumming, S. G., Burton, P. J., Joy, M., Klinkenberg, B., Schmiegelow, F. K. A., and Smith, J. N. M. (1995). Experimental habitat fragmentation and simulation of landscape dynamics in the boreal mixedwood: A pilot study. Canada-Alberta Partnership Agreement in Forestry Report A5016-137, Canadian Forest Service and the Alberta Land and Forest Service. A digital version of this report is available at `ftp://grebe.sis.ualberta.ca/pub/Papers/`.

Cumming, S. G., Burton, P. J., and Klinkenberg, B. (1996). Canadian boreal mixedwood forests may have no "representative" areas: some implications for reserve design. *Ecography*, **19**, 162–180.

DeAngelis, D. L. and Gross, L. J. (1992). *Individual-based models and approaches in ecology*. Chapman and Hall.

Demarchi, D. A. (1998). A spatial simulation model for evaluating the response of rare and endangered species to conservation strategies and forest practices: A case study on the northern spotted owl. Master's thesis, The University of British Columbia.

Eberhart, K. E. and Woodard, P. M. (1987). Distribution of residual vegetation associated with large fires in Alberta. *Can. J. For. Res.*, **17**, 1207–1212.

Lee, P. C., Crites, S., Neitfeld, M., Nguyen, H. V., and Stelfox, J. B. (1997). Characteristics and origins of deadwood material in aspen-dominated boreal forests. *Ecological Applications*, **7**, 691–701.

Lieffers, V. J., Stadt, K. J., and Navratil, S. (1996). Age structure and growth of understory white spruce under aspen. *Can. J. For. Res.*, **26**, 1002–1007.

Metz, J. A. J. and Diekmann, O. (1986). *The dynamics of physiologically structured populations*. Number 68 in Lecture Notes in Biomathematics. Springer–Verlag.

Schmiegelow, F. K. A., Machtans, C. S., and Hannon, S. J. (1997). Are boreal birds resilient to forest fragmentation? An experimental study of short-term community responses. *Ecology*, **78**, 1914–1932.

Stewart, J. D., Hogg, E. H., Hurdle, P. A., Stadt, K. J., Tollestrup, P., and Lieffers, V. J. (1998). Dispersal of white spruce seed in mature aspen stands. *Can. J. Bot.*, **76**, 181–188.