A Frame-Based Paradigm for Dynamic Ecosystem Models

A.M. Starfield
Department of Ecology, Evolution & Behavior
University of Minnesota
St. Paul, Minnesota 55108

D.H.M. Cumming and R.D. Taylor
WWF Multispecies Animal Production Systems Project
Harare, Zimbabwe

M.S. Quadling
Department of Computer Science
University of Minnesota
Minneapolis, Minnesota 55455

A Frame-Based Paradigm for Dynamic Ecosystem Models

A.M. Starfield
Department of Ecology, Evolution & Behavior
University of Minnesota
St. Paul, Minnesota 55108

D.H.M. Cumming and R.D. Taylor
WWF Multispecies Animal Production
Systems Project
Harare, Zimbabwe

M.S. Quadling
Department of Computer Science
University of Minnesota
Minneapolis, Minnesota 55455

ABSTRACT. By linking the ecological concept of a system state with the AI construct of a frame, one obtains a new paradigm for constructing models of ecosystem dynamics. The key aspect of this paradigm is that it partitions the temporal dynamics of the system; only one simple model is operational at any time. When the assumptions underlying that model are violated, a *demon* invokes rules for switching to a new frame and, hence, new model. A model to investigate interactions among rainfall, elephants, and fire in a Brachystegia woodland in Zimbabwe illustrates this approach. It shows how the paradigm provides a structure that simplifies the construction of a model. In particular, the introduction of qualitative variables, such as cool or hot fires, offers a functional approach that facilitates model construction. The example shows how a relatively simple model can lead to complex and realistic results, with unanticipated and interesting implications for management. This paradigm could lead to better communication between scientists and resource managers. It could also provide an appropriate syntax for modeling, at a certain level of resolution, the effects of climate change.

Most problems in resource management are embedded in a natural system that is driven by external forces (such as rainfall, storms, fires, sea currents, or upwelling) and characterized by internal dynamics that may include processes such as competition, succession, predation, and herbivory. On the other hand, many resource management models are concerned only with a single species such as a forest stand, an age-structured mammal population, or a fish stock.

It is not hard to find reasons why the system in which the resource is embedded is so often ignored by modelers; system models are notoriously difficult and time-consuming to build and the data to support them are seldom available. The prudent modeler may conclude that any benefit to be gained from simulating the system dynamics will be lost in the uncertainty and complexity of the exercise.

In this paper we introduce a modeling paradigm that facilitates the building of parsimonious system models. It draws on concepts and constructs from expert systems technology, but within the context of a dynamic simulation model. At the heart of the paradigm is a simple but powerful idea, one that makes connections between the ecological concept of a system state and the AI construct of a frame.

States, Frames, and Frame-Based Models

We limit our attention in this paper to a single spatially distinct region (although we suggest, at the end of the paper, how the paradigm could be extended to a patchwork of different interacting regions). If we were able to monitor the system dynamics for a very long time, we would be able to recognize distinctive *states* of the system. For example, a forest ecologist might recognize different stages of succession. A field ecologist might notice different functional attributes (e.g., nutrients are limiting in one state, while sunlight is limiting in another) and use them to define the states of the system (Noble 1987, Noble and Slatyer 1980). A theoretician might plot a trajectory of the dynamics in phase space, identify regions associated with attractors, and label these regions as the states of the system.

Modelers have previously made use of this kind of intuitive definition of system states. One example would be a markov matrix succession model (Moore 1990, Usher 1981). Another and more pertinent example is offered by Westoby et al. (1989). They use functional terms such as "grassland with many shrub seedlings" and "dense shrub cover, little grass" to characterize the states of a rangeland. They go on to develop conceptual state-and-transition models, where the transitions from one state to another are triggered by events such as "fire, following exceptional rains to provide the fuel of ephemerals."

These conceptual models offer guidance to rangeland managers, but neither markov models nor state-and-transition models attempt to simulate the underlying processes. We cannot ask these models how the system might respond to specific changes in the driving variables or to sporadic environmental events. To retain the conceptual simplicity of a state-and-transition model, but provide it with dynamic simulation capabilities, we draw on the AI concept of a *frame*.

In AI texts (such as Winston 1992) a *frame* is defined as an object possessing various types of slots or fields. Slots can be static (they contain data values), dynamic (containing procedures for computing data values or initiating actions), or they can provide connections to other frames. Frames can also have *demons*, which are procedures that watch for a specific condition and perform some action when it occurs. For example, a demon could provide missing data or update a slot under certain conditions. A frame is akin to a record, but while records are passive, a frame is designed to be both dynamic and active. A *frame system* is a network of frames connected by relations and organized into a hierarchy, with specialized slots providing the links. Frames low in the hierarchy automatically inherit properties of higher-level frames.

In our proposed modeling paradigm, we represent each state of the system by a frame. Slots within each frame contain simple mechanistic models of the key processes for the corresponding state. Demons watch for a combination of conditions that precipitate a switch from one frame to another and, when they are met, activate slots that effect the transition. Other demons may remember the values of variables prior to a frame transition, or initialize variables in the new frame.

The key components of the frame-based modeling paradigm are:

- The temporal system dynamics is partitioned into a set of states or frames. The choice of frames will depend on the purpose of the model, the way in which the system functions, or whether one set of frames rather than another facilitates the process of constructing a model.
- Independent models are constructed for each frame. These models simulate the processes that

have been identified as the key processes of interest within that frame.

- Rules are established for switching from one frame to another.

It follows that at any time during the simulation, only one model or set of simple models is operational. If the assumptions underpinning these models are violated, another model or set of models takes over.

The Construction of a Frame-Based Dynamic Model

The following is a step-by-step account of how a frame-based dynamic model is constructed:

1. As in all modeling endeavors, the first step is to identify the objectives of the model.
2. Next the driving (or input) variables are identified.
3. Then the frames are chosen. As indicated above, they could be chosen on the basis of successional stages, or recognizable alternative stable states that the system may occupy, or on a more mathematical basis with each frame representing a region of phase space around an attractor. Alternatively, the frames could be selected because they make practical sense from a management perspective, or simply because the model-builders find it useful to partition the dynamics in a certain way.
4. Next, the key variables and processes in each of the frames are identified. The objectives of the model and definition of the frames provide the structure for deciding which are the key variables and processes within each frame.
5. The pathways between frames are determined and rules constructed for switching from one frame to another.
6. Finally, a dynamic model is built for each frame. These models could be conventional numerical models or qualitative rule-based models (Starfield 1990, Starfield et al. 1989, Widman et al. 1989) or a combination of the two.

An Example: Brachystegia Woodlands in Zimbabwe

As an example, we will consider the interactions between elephants and fire in the protected areas of the Zambezi escarpments in Zimbabwe, Africa.

Background

Large areas of woodland dominated by *Brachystegia boehmii* on the Zambezi escarpment have been changed to tall grasslands or to bushlands dominated by *Combretum apiculatum* over the last three decades (Anderson and Walker 1974; Guy 1981, 1989; Taylor 1979; Thomson 1975). The changes have resulted both from increasing numbers of elephant and declining range and hence escalating densities in protected areas (Cumming 1982). Management options for retaining or recovering woodlands are control of elephant numbers and fire, but in practice these have seldom been applied in a sufficiently timely or consistent manner to achieve the desired effects (Cumming 1981a, b).

A model with multiple stable states has been built to investigate interactions between elephants and fire in East African woodlands (Dublin et al. 1990). What is needed here, however, is a model that can guide management of these ecosystems in the face of annual variability in rainfall, the hazards of uncontrolled fires, and increasing elephant populations.

Experience suggests that a Brachystegia woodland can sustain itself with small elephant populations. Larger populations of elephants, possibly aided by fire, can drive the system from woodland to a grassland interspersed with Brachystegia seedlings and shrubs. The grassland, in turn, can regenerate the woodland, provided it is protected from hot fires and elephant densities are low. However, fire and larger elephant densities can suppress Brachystegia shrubs and eventually drive the grassland to a bushland where Combretum and other fire-resistant shrubs compete with the Brachystegia shrubs. Once the fire-resistant shrubs have become established, a return to grassland is prevented, although over a long period (of about 50 years) Brachystegia shrubs might push through the cover to dominate the canopy and return the system to a Brachystegia woodland again.

The Model

The objectives of the model are:

- To draw on the available ecological knowledge and/or data in order to compare the effects of various management strategies on a time-scale of 25 to 100 years.
- To explore the interaction between ongoing processes and events that drive the system, again from a management perspective.

The time step chosen for this model is one year. The driving variables are:

elephants: a qualitative variable with levels 0, 1, 2, 3, or 4 (corresponding to no elephants, densities of about 0.25 per sq km, 0.50 per sq km, 0.75 per sq km, and 1.00 per sq km or higher);

rainfall: annual rainfall with two states, *low* and *high* ;

match: can be *on* or *off*, representing the inception of a fire (either naturally or deliberately);

burntime: can be *early* or *late*, representing the time (during the dry season) when a fire occurs.

The driving variable *match* implies that an attempt is made to ignite a fire, but it does not necessarily follow that there is sufficient fuel to sustain it. We therefore introduce a variable *fuel-load* on a scale of 0 to 6 (or *very low* to *very high*). Depending on *fuel-load* (and also on whether the fire occurs in the early or late dry season), the fire will be either *cool* or *hot*. We introduce the qualitative variable *fire type* with three states (0 = *no fire*, 1 = *cool fire*, and 2 = *hot fire*). Although these variables are common to all frames, the rules that govern them are not; fuel accumulates and fires develop in different ways in the different frames.

Based on current understanding of the long-term dynamics, it was decided to introduce three frames: Woodland, Grassland and Bushland. The Woodland frame represents a mixed woodland dominated by *Brachystegia boehmii*. The Grassland frame represents a tall grassland with occasional Brachystegia trees and a shrub layer dominated by Brachystegia shrubs together with Combretum and other species of shrubs and emergents. The Bushland frame represents a mixed scrubland growing through to wooded bushland with *Combretum apiculatum* and other Combretum species dominating.

The Bushland frame. It is easiest to describe the model working backwards from the simplest frame (Bushland) to the most complex (Woodland). Since the time scale of our objectives is 25 to 100 years, we chose to ignore the possible 50-year transition from Bushland to Woodland. It is possible for the system to switch back to a grassland, but this can only occur during the first five or six years when Combretum shrubs are still vulnerable to fire; otherwise the system is trapped in the Bushland frame for the rest of the simulation (i.e., Bushland is an absorbing state).

A switch back to Grassland may be triggered by a hot fire during the first six years in the Bushland frame. We therefore introduce a variable *bushtime* to keep track of the number of years since the switch from Grassland to Bushland.

While *bushtime* is less than 7, the following rules apply:

1. The variable *fuel-load* is reset to zero after any fire.
2. *Fuel-load* is increased by one if *rainfall* is *low*, and by two if *rainfall* is *high*.
3. A fire can only occur if *match* is *on* (ignition) and the fuel-load is greater than or equal to 1. The fire will always be *cool* if the ignition is early or if the fuel-load is 1. If ignition is late, the fire will be *hot* if the fuel-load is 3 or greater, and can be either hot or cool (with equal probability) if the fuel-load is 2.

The rule for switching from Bushland to Grassland is:

Switch definitely if there is a hot fire and *bushtime* < 3, with probability 0.75 if *bushtime* is 3 or 4, and with probability 0.5 if *bushtime* is 5 or 6. Otherwise the system is trapped in the Bushland frame for the rest of the simulation. What occurs in the Bushland frame during that time is irrelevant.

The Grassland frame. In the Grassland frame we represent the Brachystegia shrubs by the variable *height* on a scale of 0 to 3. We also differentiate between the height category which is an integer in the range 0 to 3, and the height number which is a real number between 0 and 3.99. This distinction permits a mix of numerical and qualitative modeling with advantages that will soon become apparent.

Rules 1, 2, and 3 from the Bushland frame are repeated in the Grassland frame; i.e., there is no change in the rules governing fuel-load and fire. In addition, there is a set of rules for increasing or decreasing the height of Brachystegia shrubs (a demon sets *height* equal to 1 whenever the system switches into the Grassland frame):

1. Low rainfall adds 0.2 to the height number, while high rainfall adds 0.25.
2. The following table determines how much to subtract from the height number, depending on the category of elephants and on the height category. (The height category is just the truncated integer value of the height variable.)

		Brachystegia height category		
		1	2	3
	1	0.075	0.05	0.03
Elephant	2	0.15	0.10	0.06
density	3	0.25	0.20	0.10
	4	0.35	0.35	0.35

(Small Brachystegia bushes which are in the field or grass layer, i.e., in height category 0, do not seem to be eaten by elephants.)

3. Shrub height can also be reduced by hot fires. If the height category is 0, 1, or 2, then we subtract 1.0 from the height number if the fuel-load is less than 3; for higher fuel-loads we subtract 2.0. Taller shrubs (height category 3) will only sometimes be affected by hot fires; the probability of this happening is specified as a function of fuel-load in a table, and we subtract 2.0 whenever a random number is less than the appropriate table entry.

(Notice how in the above rules we have exploited the dual characterization of the height variable. It is convenient to think of height as a real variable in rule 1 and in the table entries of rule 2. This allows us to combine fractional increases and decreases in height as a function of rainfall and herbivory. On the other hand, it is equally convenient to use the height category to specify the columns in the look-up table in rule 2).

There are three conditions that will trigger a switch from Grassland to Woodland:

1. If the shrub height has been continuously in category 3 (the highest category) for five years.
2. If the annual rainfall for the current year is high and the shrub height has been continuously in category 3 for at least three years.
3. Irrespective of the rainfall, with probability 0.5 if the height has been in category 3 for three years and probability 0.75 in the fourth year.

The switch from Grassland to Bushland is associated with a depletion of *Brachystegia boehmii* seedlings and shrubs and the encroachment of other shrubs and emergent tree species. The rule is:

Switch if the height variable has been continuously less than 1 for at least 5 years and there have been no hot fires for the past two years.

The Woodland frame. Here we track the state of the mature Brachystegia trees using a canopy variable on a logarithmic scale as follows:

Canopy category	Representing (% canopy)
1	0 to 4
2	4 to 12
3	12 to 28
4	28 to 60
5	> 60

The *canopy* variable can be specified as a category (integer) or number (real variable) as with the bush height variable.

The fuel-load model is slightly more complex here than in the other two frames, since fuel-load builds up more slowly as the canopy increases. Rule 2 of the Bushland frame is replaced by a table specifying the annual increment in fuel-load as a function of the rainfall and the canopy category. Rules 1 and 3 from the Bushland frame are copied without change.

The canopy number is randomly initialized to a value between 3.5 and 4.5 whenever there is a switch from the Grassland frame. The following rules then apply to changes in canopy:

1. Add 0.11 to the canopy number in a low rainfall year and 0.22 in a high rainfall year.
2. The amount to be subtracted from the canopy number each year because of elephants' herbivory is determined from the following table:

		Canopy category				
		1	2	3	4	5
Elephant density	1	0.16	0.14	0.12	0.10	0.075
	2	0.35	0.30	0.25	0.20	0.15
	3	0.70	0.60	0.50	0.50	0.50
	4	1.00	0.80	0.80	0.80	0.80

3. Subtract an amount from the canopy number if there is a hot fire. This amount is specified as a function of the fuel-load in a table.

The switch from the Woodland frame to the Grassland frame occurs when mature trees have been removed from the canopy. The rule is:

Switch definitely in a year in which the canopy category is 1 and with probability 0.3 if the category is 2.

Figure 1 summarizes the transitions between the frames.

A Frame-Based Implementation

Figure 2 demonstrates how the frame representation can be used to implement the Brachystegia model. In addition to the frames corresponding to each state of the system, we introduce a global frame which describes the processes common to each state and determines the transitions between states.

Each frame contains a name slot for identification, a description slot containing a short verbal description of the processes modeled in that frame, and a list of variables. In addition, a parent slot is used to establish whether a frame is a specialization (child) of another frame. Thus the Bushland, Grassland, and Woodland frames are all child frames of the global frame and consequently inherit the models and variables contained in the global frame.

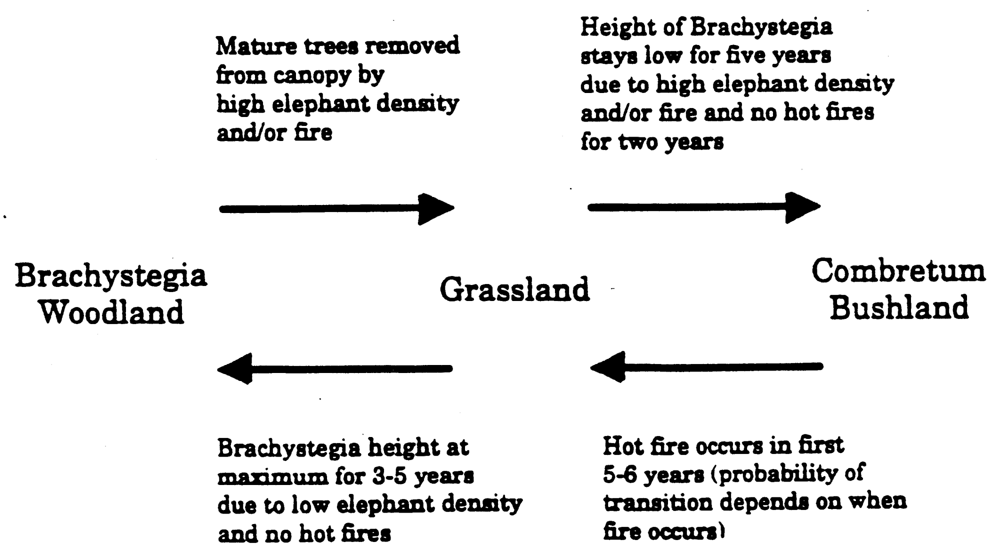

Figure 1. *A summary of the transition rules in the Brachystegia model.*

Figure 2. A frame-based representation of the Brachystegia model.

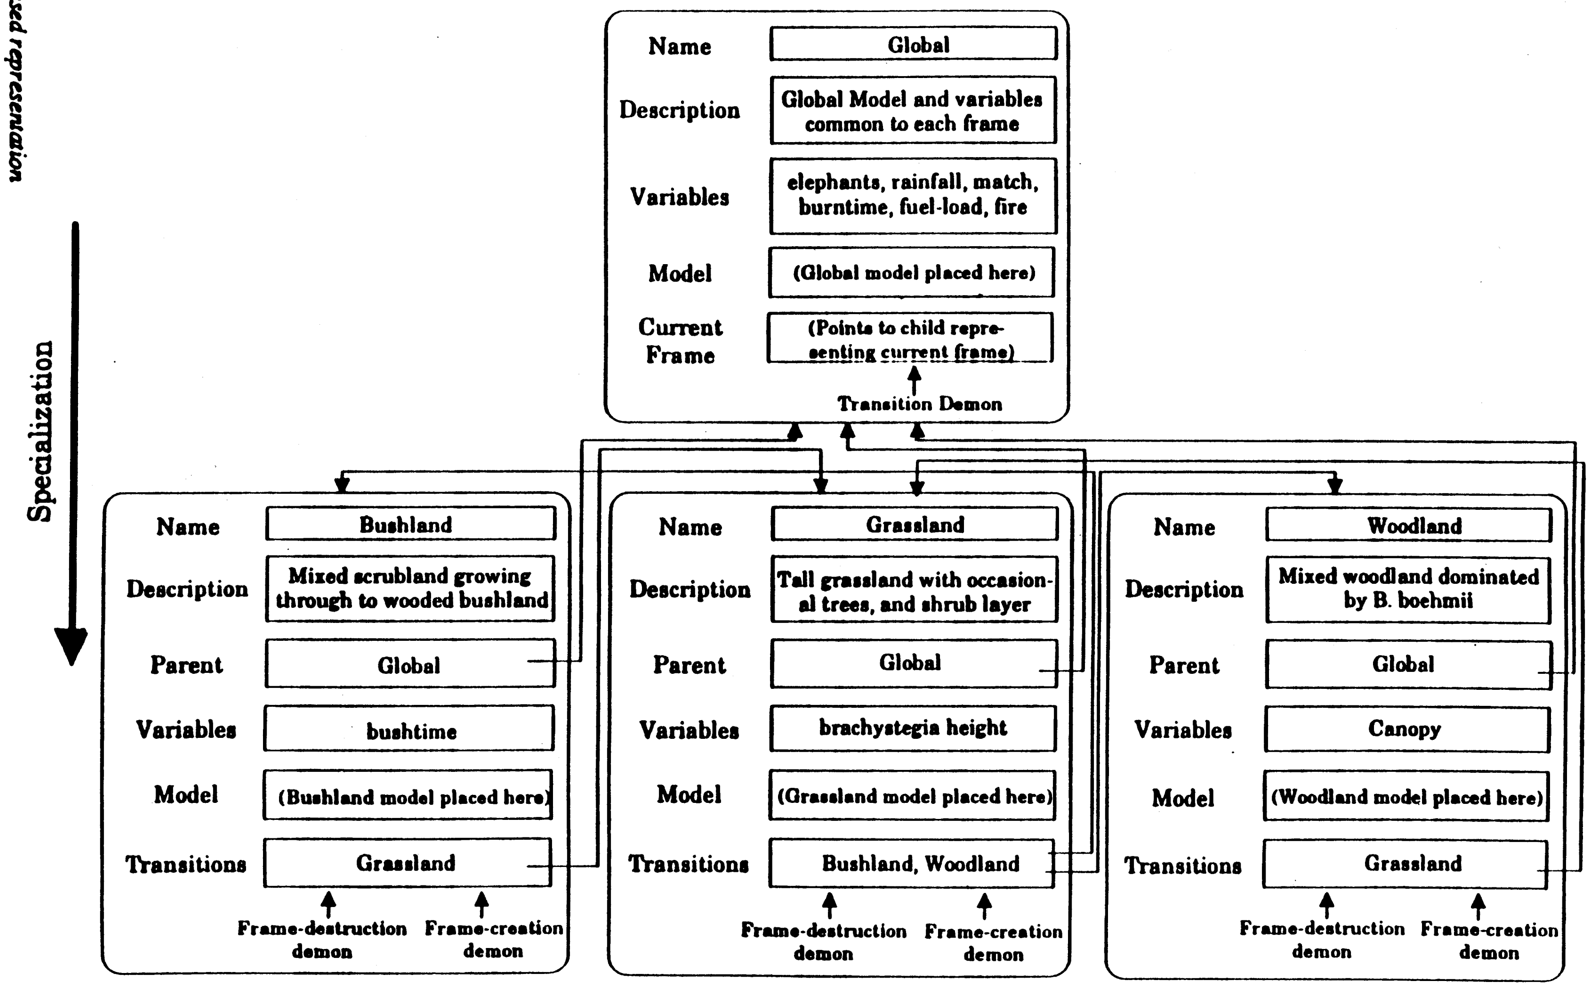

The global frame contains a *current frame* slot which is used during the dynamic simulation to indicate the active frame (or current state of the system). A transition demon watches for conditions under which a new frame would become the current frame. When a transition is triggered, a *frame destruction* demon is activated in the current child frame and a *frame creation* demon is activated in the new child frame. These demons save information on the frame that has been vacated and initialize variables in the new frame.

Results

The model can be used in different modes. For example, one can perform individual simulations with all driving variables input, year by year, by the user. In this way the user can detect possible 'bugs' in the code or in the model itself, fine tune some of the parameters of the model, look for unrealistic output, make comparisons with field data (where these are available) and ultimately gain confidence in the model.

As a next step, the driving variables can be generated from probability distributions built into the model. In this mode one can also test and gain confidence in the model by trying to interpret the output. Figure 3 is an example of output obtained in this way. It shows a Woodland switching to Grassland, the Grassland switching to Bushland, recovering as the result of a hot fire, and then switching irrevocably back to Bushland. (Notice, incidentally, how the variable *canopy* is only defined when the model is in frame 1, the Woodland frame. Similarly *height* is only defined in frame 2, the Grassland frame.)

It should be stressed that Figure 3 was gen-

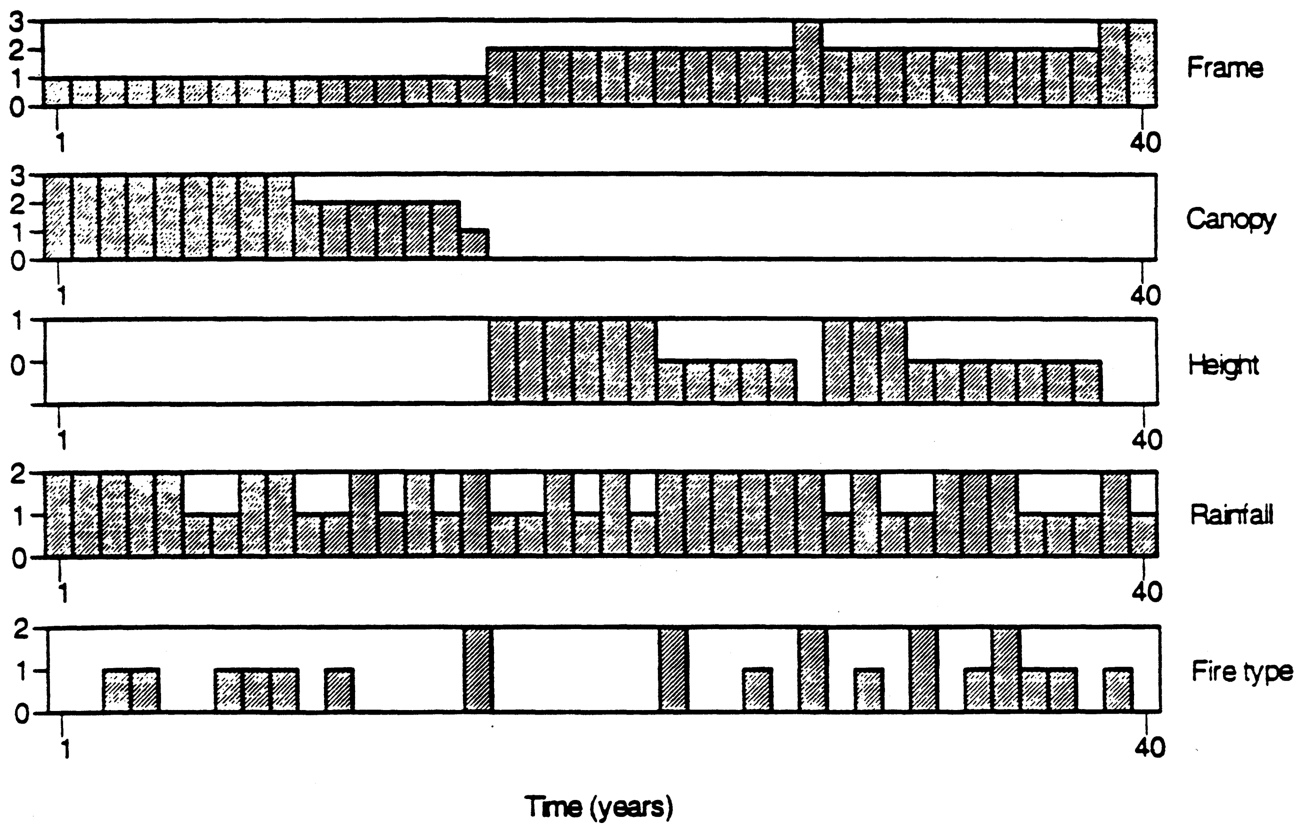

Figure 3. *An example of output from a single replicate of the Brachystegia model, showing switches from Woodland (frame 1) to Grassland (frame 2), Bushland (frame 3), back to Grassland and then back to Bushland. Level 1 corresponds to low annual rainfall, 2 to high annual rainfall. The Fire variable is 0 for no fire, 1 for a cool fire, and 2 for a hot fire.*

erated by one replicate only. Other replicates, using exactly the same probability distributions, showed a variety of behaviors, depending on rainfall patterns, when fires occurred, and whether fires were early or late, cool or hot. It follows that if one is trying to detect trends or wants to compare strategies or perform sensitivity analyses, it is essential to look at the output from a large number of replicates. The management strategies, as well as the driving variables, are hard-wired into the model in this mode.

Figure 4 is an example of output one can obtain from a series of replicated computer experiments. The figure plots the probability of ending a 100-year simulation in the Bushland frame versus the probability of igniting a fire. Results are shown for four different scenarios. The figure illustrates how the probability of being trapped in the Bushland frame increases (at low ignition probabilities) with an increasing probability of ignition, which corresponds with experience. It also shows how, at high ignition probabilities, the likelihood of ending in the Bushland frame is reduced substantially as the probability of early (and hence cool) fires is increased. This too corresponds with experience.

The exception to these trends is the case where the elephant category is set at two, suggesting a new hypothesis: the system is less sensitive to the frequency and timing of fires when elephant densities remain above a certain threshold.

Figure 5 is also an example of output from a series of replicated computer experiments, designed in this case to explore sensitivity to the probability of high rainfall. It could equally well be interpreted as an experiment to explore the likely consequences of climate change. It

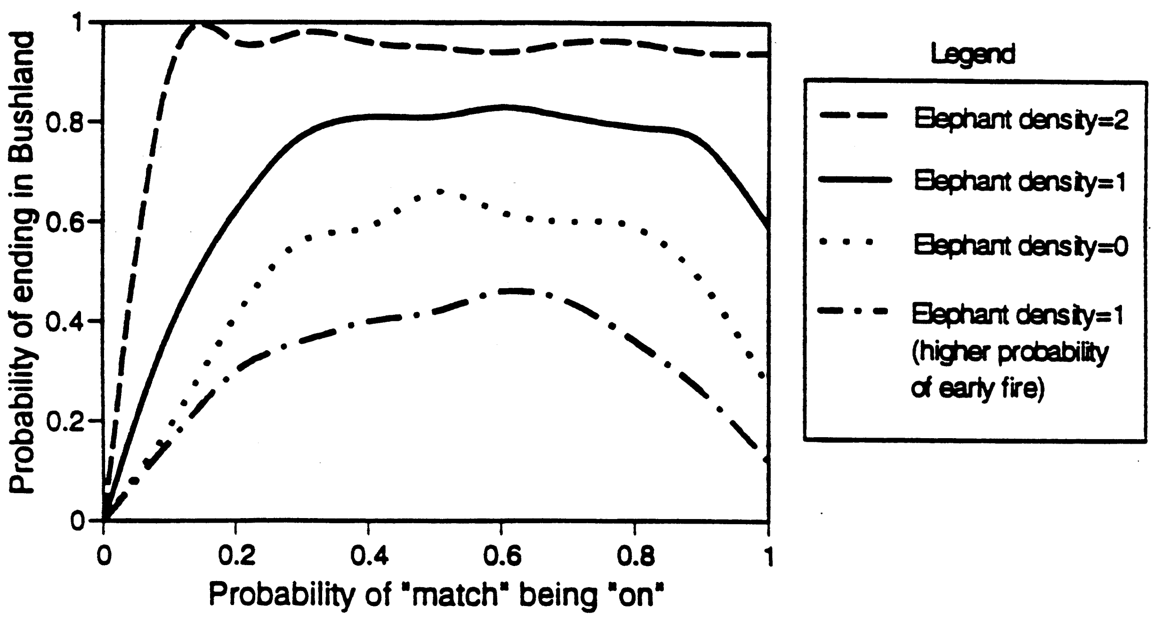

Figure 4. *A summary of a number of computer experiments, all starting in the Woodland frame and running for 100 years. The results show the probability of an undesirable outcome—ending the simulation in the Bushland frame—and how it depends on elephant densities and fire frequencies. Each point on a graph was obtained from 500 replicates. In all cases the probability of high rainfall was 0.5. The probability of an early fire was 0.3, except for one experiment where that probability was increased to 0.7.*

suggests that frequent burning would no longer be a viable management option if the frequency of high rainfall years were to increase.

Discussion

The purpose of this paper is to introduce a modeling paradigm and give some idea of how it might be used. There are a number of points to be made about the approach, the way in which it might be implemented, and its implications in resource management.

The novelty of this paradigm lies first in the idea of constructing separate models for each frame (or state) of the system, and second in using rules to decide whether conditions have changed so as to trigger a switch from one frame to another. At any time only one simple model is operating. It is totally independent of all the other models and does not call on them or interact with them in any way. When the assumptions behind the operating model are no longer applicable, it is retired and another model takes its place.

A consequence of this approach is a huge simplification in the task of constructing a system model. Each time a new frame is introduced, so too are the assumptions that define it. These assumptions, in turn, help to identify the key processes and variables. Arguments about what to include or leave out are easily resolved by referring to the frame definition. If necessary, the definition can be expanded or additional frames can be introduced, but at all times the to-and-fro discussion that is so important a part of model building is guided, unambiguously, by the frame structure.

Useful models can be built and tested quickly. (The Brachystegia Woodland example described

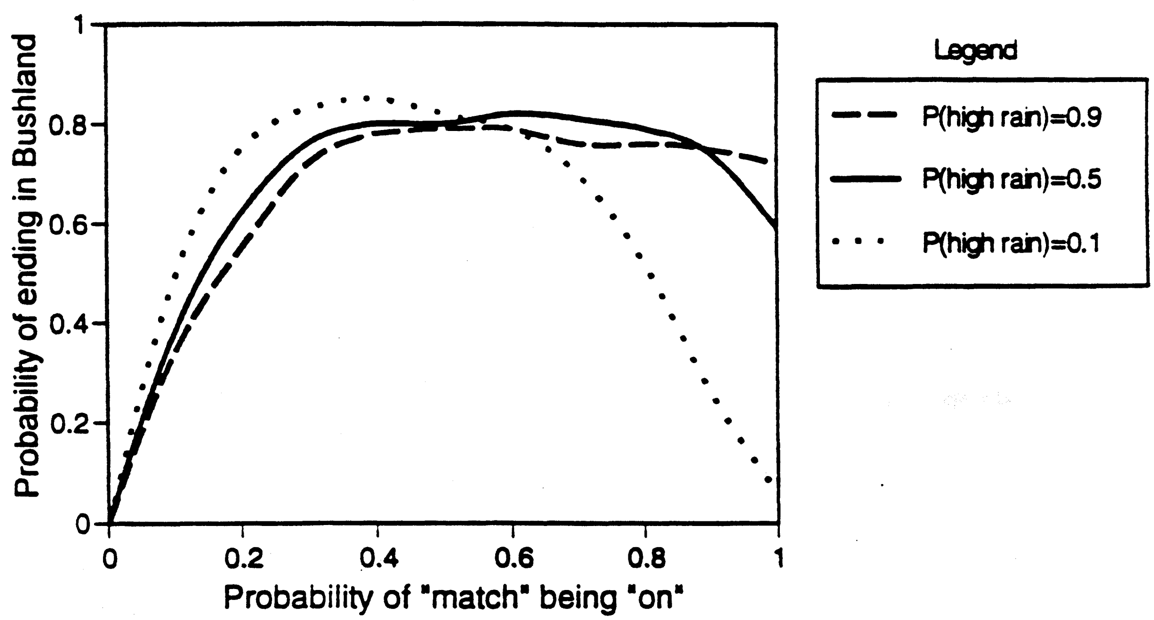

Figure 5. *An experiment to explore possible effects of climate change. It suggests that a strategy of frequent burning is only effective if there is a low probability of high rainfall. In these simulations the elephant density was at level 1 and the probability of an early fire was 0.3.*

above, was constructed and tested by a team of three working for two days.) The approach lends itself to rapid prototyping; ideas can be tested and rejected without tearing the whole fabric of a large and complicated model. All too often, system models get bogged down in details and are redesigned before they have been completed and tested. Frame-based modeling encourages modelers first to build and test simple models and then, guided by that experience, to add detail (where necessary) as part of a sensitivity analysis of the working model.

It is not obligatory to use qualitative variables in a frame-based model, but qualitative variables defined in a functional manner also facilitate the process of constructing a model. For example, the definition of a hot fire in the Brachystegia model prescribes its effect. Similarly, the definition of shrub height in the Grassland frame simplifies the herbivory rules; by definition, elephants do not eat shrubs of height category zero in the grass layer.

The combination of frames and qualitative variables not only makes it easier to construct models, it also makes it easier to describe them. The result is vastly improved communication between modelers and managers.

Managers and scientists alike tend to develop hypotheses or heuristics based on their own experience, but they seldom know whether what they have observed, over periods of even 20 or 30 years, is truly representational or just a sequence of perhaps unlikely events. For example, a management heuristic that has been recognized is "set frequent early fires to promote Woodland." Figures 4 and 5 suggest that this will be an effective management strategy only if it is applied rigorously (the probability of an early fire must be very high) and then only under certain conditions; elephant densities cannot be too high and the probability of high rainfall must be below a threshold.

Figures 3, 4 and 5 demonstrate that even a relatively small and simple model can lead to intriguing results. A model such as the Brachystegia model provides just enough complexity to put experience into a more realistic context. We would argue that the usefulness of the model is a consequence of the modeling paradigm; the combination of frames and qualitative variables lends itself to exploring interactions between management, ongoing processes (such as herbivory) and influential events (such as fires) at a level of resolution that is simple enough to understand but complex enough to give interesting and unanticipated results.

For the same reasons, we believe this is an appropriate paradigm (as Figure 5 suggests) for exploring the likely effects of long-term climate change. Climate change leads to shifts in the frequency of rainfall, or shifts from early to late rains, or an increased probability of storms, or similar types of effects. All of these can be captured effectively but parsimoniously in a qualitative frame-based model.

Some of the above points apply to good modeling practice in general and are not an exclusive consequence of frame-based modeling. However, frame-based modeling makes it easier to implement good modeling practices in the same way as proponents of certain computer languages would argue that their syntax encourages good programming practices.

All of the above points are concerned with the model itself and are independent of how the model is actually implemented on a computer. Although this modeling paradigm was inspired by the AI concepts of frames, slots, and demons, these formal constructs are not actually necessary for implementing *small* dynamic models such as the Brachystegia Woodland model. (It was, in fact, written in Pascal.) However, a formal frame structure, such as that depicted in Figure 2, offers a number of benefits:

- It has been designed to closely resemble the overall structure of the systems being modeled. The immediate benefit of this is that the modeler is aided considerably in describing these systems.
- It lends itself to the design of user-friendly software. With a few small changes the structure depicted in Figure 2 could be used in a graphical user Interface to assist the modeler. This kind of structure would be essential if the model has many frames and interactions.
- The concept of inheritance in frame systems facilitates the implementation of hierarchical frame models and ties in well with hierarchical

descriptions of ecosystems (O'Neill et al. 1986).

- Just as some expert systems provide explanations, so one could imagine extending the frame-based representation to include a run-time explanation feature.

For example, some of the results in Figure 4 could be explained as follows: At low ignition probabilities, Brachystegia shrubs continue to grow despite elephant herbivory. As the ignition probability increases, so it becomes more likely that the combination of fire and elephant herbivory will suppress Brachystegia shrubs and facilitate a switch from Grassland to Bushland. At very high ignition probabilities, there is an increased likelihood of a hot fire in the first few years after a switch to Bushland, and so there is an increased probability of a switch back from Bushland to Grassland.

It is unlikely that the computer could provide an explanation of this kind, but it would be an interesting research problem to design a feature that could provide snippets of explanation which would enable the user to construct a full explanation.

- Finally, one can also imagine extending the paradigm to a spatial patchwork of interacting regions. For each region, the temporal dynamics would be simulated by a frame-based model. In some cases the same model might apply to all regions, although different regions at a particular time might be found in different frames. In other cases, there could be different models (and hence sets of frames) associated with different regions. The frame models and frame transition rules would be extended to include interactions between nearest neighbors.

Cellular automata "are systems of cells interacting in a simple way but displaying complex overall behavior" (Wolfram 1985). Each cell may be in one of a finite number of possible states, and at each time step the cells update their states according to a transition rule which depends on the states of the neighboring cells. Cellular automata models are having an impact on our understanding of how ecological processes engender and maintain patchiness and how we interpret spatial patterns (Phipps 1992). If we were to add a spatial component to frame-based models, we would have the spatial power of cellular automata models with the temporal power of the frame-based paradigm. This could provide the appropriate level of resolution (neither too detailed nor too simple) for models addressing regional problems.

Acknowledgments

The senior author gratefully acknowledges support from the WWF Multispecies Animal Production Systems Project, the South African Nature Foundation, and the Office of International Education at the University of Minnesota. We are grateful to the Zimbabwe Department of National Parks and Wildlife Management and to Ian Coulson for facilities provided at the Sengwa Wildlife Research Institute during the course of this work. Finally, we appreciate the comments of those who reviewed the first draft of this manuscript and one (anonymous) reviewer in particular should notice we have made liberal use of his or her suggestions.

References

Anderson, D.G., and B.H. Walker. 1974. Vegetation composition and elephant damage in the Sengwa Wildlife Research Area, Rhodesia. Journal of the South African Wildlife Management Association 4: 1-14.

Cumming, D.H.M. 1981a. The management of elephant and other large herbivores in Zimbabwe. Pages 91-118 in: Problems in Management of Locally Abundant Wild Mammals, P.A. Jewell, S. Holt, and D. Hart, editors. Academic Press, New York.

Cumming, D.H.M. 1981b. Elephant and woodlands in Chizarira National Park, Zimbabwe. Pages 347-349 in: Problems in Management of Locally Abundant Wild Mammals, P.A. Jewell, S. Holt, and D. Hart, editors. Academic Press, New York.

Cumming, D.H.M. 1982. The influence of large herbivores on savanna structure in Africa. Pages 217-245 in: Ecology of Tropical Savannas, B.J. Huntley and B.H. Walker, editors. Springer Verlag, Berlin.

Dublin, H.T., A.R.E. Sinclair, and J. McGlade. 1990. Elephants and fire as causes of multiple stable states in the Serengeti-Mara woodlands. Journal of Animal Ecology 59: 1147-1164.

Guy, P.R. 1981. Changes in the biomass and productivity of woodlands in the Sengwa Wildlife Research Area, Zimbabwe. Journal of Applied Ecology 18: 507-519.

Guy, P.R. 1989. The influence of elephants and fire on a *Brachystegia-Julbernadia* woodland in Zimbabwe. Journal of Tropical Ecology 5: 215-226.

Moore, A.D. 1990. The semi-Markov process: a useful tool in the analysis of vegetation dynamics for management. Journal of Environmental Management 30: 111-130.

Noble, I.R. 1987. The role of expert systems in vegetation science. Vegetatio 69: 115-121.

Noble, I.R., and R.O. Slatyer. 1980. The use of vital attributes to predict successional changes in plant communities subject to recurrent disturbances. Vegetatio 43: 5-21.

O'Neill, R.V., D.L. DeAngelis, J.B. Waide, and T.F.H. Allen. 1986. A Hierarchical Concept of Ecosystems. Princeton University Press, Princeton, New Jersey.

Phipps, M.J. 1992. From local to global: the lesson of cellular automata. Pages 165-187 in: Individual-based Models and Approaches in Ecology, D.L. DeAngelis and L.J. Gross, editors. Chapman & Hall, New York.

Starfield, A.M. 1990. Qualitative, rule-based modeling. BioScience 40: 601-604.

Starfield, A.M., B.P. Farm, and R.H. Taylor. 1989. A rule-based ecological model for the management of an estuarine lake. Ecological Modelling 46: 107-119.

Taylor, R.D. 1979. Biological conservation in Matusadona National Park, Kariba. Transactions of the Rhodesia Scientific Association 59: 30-40.

Thomson, P.J. 1975. The role of elephants, fire and other agents in the decline of *Brachystegia boehmii* woodland. Journal of the Southern African Wildlife Management Association 5: 11-18.

Usher, M.B. 1981. Modelling ecological succession, with particular reference to Markovian models. Vegetatio 46: 11-18.

Westoby, M., B. Walker, and I. Noy-Meir. 1989. Opportunistic management for rangelands not at equilibrium. Journal of Range Management 42: 266-274.

Widman, L.E., K.A. Loparo, and N.R. Nielsen, editors. 1989. Artificial Intelligence, Simulation and Modeling. Wiley, New York.

Winston, P.H. 1992. Artificial Intelligence. Third edition. Addison-Wesley, Reading, Massachusetts.

Wolfram, S. 1985. Some recent results and questions about cellular automata. Pages 153-167 in: Dynamical Systems and Cellular Automata, J. Demongot, E. Goles and M. Tchuente, editors. Academic Press, London.

Anthony Starfield was educated as an applied mathematician at the University of the Witwatersrand in Johannesburg, South Africa. He has worked and published in such diverse areas as numerical analysis, mine ventilation, geomechanics, engineering education, expert systems, and ecological modeling. The unifying theme of his interests is modeling, the methodology of modeling, and how to teach modeling. He has coauthored four books, including *Building Models for Conservation and Wildlife Management* (second edition, Bellwether Press, Edina, Minnesota) and *How to Model It: Problem-Solving for the Computer Age* (McGraw-Hill). He is currently a professor in the Department of Ecology, Evolution & Behavior at the University of Minnesota and is actively involved in the Conservation Biology program.

David Cumming, a graduate of Rhodes University, South Africa, has been involved in wildlife research and conservation in Zimbabwe for 28 years. He was Chief Ecologist and then Deputy Director of the Zimbabwe Department of National Parks and Wildlife Management before becoming Project Leader of the WWF Multispecies Animal Productions System Project in 1988. This project is examining ecological and economic features of single and multispecies animal production systems in African savannas and assists in the implementation of Zimbabwe's Communal Areas Management Programme for Indigenous Resources (CAMPFIRE). His current research and development interests center on natural resource management in African rangelands and particularly the role of multispecies systems in sustainable land use.

Russell Taylor, a graduate of the University of Zimbabwe, has been involved in wildlife management and conservation in Zimbabwe for the past 20 years. His research has focused on large herbivores and herbivore-plant interactions, particularly of buffalo and elephant. He was a Principal Ecologist in Zimbabwe's Department of National Parks and Wildlife Management before taking up the position of Ecologist on the WWF Multispecies Animal Productions System Project. His major current activity is in the field of community-based resource management and utilization.

Mark Quadling studied physics and applied mathematics at the University of the Witwatersrand, Johannesburg, South Africa. He is currently in the Ph.D. program in Computer Science at the University of Minnesota.