

Dynamics Modelling and Culling

Stephen Chenney Jeffrey Ichnowski
David Forsyth
University of California at Berkeley

November 10, 1998

Abstract

Procedural animation, via dynamical systems, has many advantages over keyframe animation, yet suffers from high computational cost and difficulties in modeling. We describe tools that analyze and process systems to enable culling if the system is not in view, based on a qualitative model of viewer prediction. Application of these tools results in the cost of dynamic simulation growing linearly with the number of systems in view, regardless of the number of systems out of view. To address the difficulties of modeling, we present a runtime layer that interfaces between a VRML browser and simple descriptions of dynamics written in Java. To create these runtime environments, we have developed a modeling program in which a user interactively builds models and previews their dynamics. Models created using these tools are presented, and their performance demonstrated. An important outcome of this work is the emergence of a library of dynamical systems, each of which is simple to incorporate into a virtual world. Together, our tools enable large numbers of complex dynamic models to be efficiently and easily included in a VRML world while maintaining high frame rates.

1 Introduction

Animating geometry procedurally, using a dynamical system¹ rather than with keyframes, offers the advantages of physical realism, interactivity, compact descriptions and infinite variety of motion. However, the use of dynamical systems has been inhibited by two factors: their computational cost, even when out of view, and the difficulties of implementation, partially due to a lack of standard modeling and runtime environments.

VRML is particularly prone to these problems. The most suitable, currently available, scripting language is Java. Controlling a VRML world from a Java object requires knowledge of a large number of methods and types, and is highly

¹A dynamical system consists of a set of state variables and a set of equations for describing how those variables change over time.

error prone. A lack of debugging environments, or even consistent exception handling, only exacerbates the problem.

The aim of our research is to develop modeling tools that enable large numbers of efficient dynamical systems to be incorporated into virtual environments, while abstracting the modeling process as much as possible. To achieve efficiency, we concentrate on *culling* dynamical systems: if the system is not in view, we do not compute any dynamics for it.

As a concrete example of the benefits of culling, consider a museum of kinetic sculpture (such as mobiles). In such a world there are a large number of dynamic objects – the exhibits – placed in mutually non-visible rooms. Traditional dynamic simulation requires computing the motion of all the sculptures on every frame, even if some sculptures are not visible to the viewer. Naturally, this would severely limit the size of the museum. Culling avoids dynamic state computation for the exhibits not in view, just as geometry culling avoids rendering their geometry. Exploiting culling allows the museum to be arbitrarily large, provided only a few rooms are visible at once. There are similar examples in factory simulations or dense city environments.

In this article we describe three tools that together provide an environment for authoring cullable dynamic rigid-body objects in VRML and Java:

- A code transformation tool that exploits **approximations to dynamical systems** to enable culling. The frame rate for a world using these systems depends only on the number of systems in view. This results in significant speedups over conventional models and enables very large virtual worlds.
- A **runtime layer** generator, which defines a simple standard interface between a VRML browser and dynamical systems described in Java. This hides from authors the complexities of the underlying interface, while also allowing for a library of commonly used dynamical systems, both of which reduce the difficulties of generating dynamic content.
- A **rigid-body modeler** which allows users to interactively design the runtime layer and preview the dynamic behavior. The modeler environment is written in Java, so some degree of interactive debugging is possible.

This article describes these tools, including some example systems and a discussion of the runtime performance improvements obtained. Our tools are applicable if the spatial range of the dynamic model may be bounded by a static volume, if the model is closed to outside influence, if the underlying equations are continuous, and if the dimension (number of degrees of freedom) of the system is small. Note that, while this article focuses on VRML and Java as the target environment, the underlying techniques are applicable to any rendering and language environment.

2 Culling and Prediction

In previous work [2], we have described the use of approximations to cull moving objects. Those techniques employed statistical models and neural networks in hand-coded systems to enable the culling of dynamics, allowing new dynamic state to be generated rapidly even after long periods out of view. In this article we describe automated tools to perform the same task for systems whose range of motions can be bounded tightly. Where such bounds exist, the key problem — the **consistency** problem — is this²:

When a viewer turns away from a dynamic object, which is culled, and then turns back to the object, the new dynamic state should be consistent with a viewer’s reasonable predictions for the behavior of the object during the time it was out of view. A solution to this problem requires the generation of new dynamic state over very long time-steps (1) within an error bound determined by a viewer’s ability to predict and (2) without significantly delaying rendering of the next frame.

It helps to consider the viewer as an adversary who is always seeking to discover a contradiction in state as objects move in and out of view. It is the system’s job to make sure that the viewer never spots such an inconsistency and never experiences significant lag.

The naive way to generate consistent state is to determine exactly how the system behaved while it was out of view, and display the resulting state. This satisfies (1) above, but the resulting state may take a very long time to compute — arbitrarily long if an object can be out of view for arbitrary periods of time — and the method fails on point (2) above.

A better way of generating new state when an object re-enters the view comes from observing that viewers cannot accurately predict the behavior of a system over time [3]. If accurate predictions were possible, a viewer could know all the future of a world simply by observing it for a short period of time, which is obviously not the case. Solving consistency means satisfying viewer predictions, so if viewers cannot predict everything, we need not compute everything. All we need do is compute those things a viewer can predict. This is a fairly flexible requirement — we are free to do more or less work depending on the quality of the simulation required. For entertainment purposes, larger errors in state may be acceptable, whereas visualization tasks or training simulations generally require higher accuracy.

A viewer’s inability to predict is the result of several factors: uncertainty in the last known state of the system; uncertainty in the details of the model; and a lack of knowledge about factors that influence the system. Our methods are based primarily on uncertainty in the last known state of the system. In other

²Previously work [2] defines two other problems: completeness and causality. These problems are not relevant if the objects can be tightly bound, and will not be considered in this article.

words, when a system leaves the view, it is not possible for a viewer to have a completely accurate picture of, for instance, how fast it’s moving, because of inaccuracies in the rendering (timing inconsistencies and pixel sampling errors) and limitations in the viewer’s perception. By taking the uncertainty, and propagating it forward through the time an object is out of view, it is possible to find out how uncertain an object’s state is when it re-enters the view. Provided the new state we generate is within this re-entry region of uncertainty, a viewer cannot detect an inconsistency.

Our methods for generating new state will rely heavily on the approximation of dynamical systems. Approximations have previously been exploited by several authors for a variety of purposes. Hodgins and Carlson [6] and Setas et. al. [8] employ approximations for dynamical systems that are further from the viewer — level of detail for dynamic simulations. Grzeszczuk et. al. [4] employ neural networks as approximation functions for dynamical system, thus reducing the cost of generating state for each frame, but ignoring issues such as level of detail and culling.

3 Generating Approximations

Our tools takes as input a basic description of the dynamics and produce an alternate description suitable for culling. We assume the input description is the most accurate required for simulation, and we will refer to it as the *accurate model*. We will call the output description the *efficient model*. The tools we describe work for dynamical systems that are free from external influence, and have a state space of low dimension. They are only useful if the accurate model is expensive to evaluate over long time intervals. This is the case with, for instance, systems of differential equations evaluated by numerical integration, or systems described by state machine transitions.

In our implementation, the accurate model consists of a Java class file which implements a function for evaluating the system at some given time when also given an initial state and time. The function should be able to generate state in real time (it will be used when the system is in view) but need not be significantly faster. In specifying the accurate model, the user also defines other information helpful in analysis, such as initial conditions (which may be fixed or random), the dimension of the system, and a number of other parameters.

We distinguish systems according to whether they are periodic or not. A system is periodic if its state at time t , S_t , is identical to that at time $t + T$, where T is the period of the system. If $S_t = S_{t+T}$, then $S_t = S_{t+nT}$ for any integer n . If no such T exists, then the system is not periodic.

3.1 Periodic Systems

As an illustrative example of a periodic system, consider a roller coaster model. The car runs on a track described by a uniform cubic b-spline (figure 1), under the influence of gravity but without friction. The “upright” direction of the car




Figure 1: *The roller coaster model.*

is also described with a b-spline. Two state variables describe its motion: the parametric position on the track, u , and the derivative of u with respect to time, \dot{u} .

Consider what a viewer might be able to predict about a system like the roller coaster. Over very short periods, on the order of a couple of seconds, the smooth nature of the motion makes accurate prediction easy: a viewer can simply extrapolate the last seen position and velocity. But if the roller coaster has been out of view for longer, the viewer must rely on what they know about roller coasters in general, such as how fast the cars move at a given position on the track. Predictions of this type will generally be uncertain: as long as the car appears to be doing about the right thing, the viewer will think everything is reasonable. Our task is to find functions that are quick to evaluate and whose output is close enough to the true dynamics to fool the viewer.

The behavior of a periodic system can be completely described by its state function over the course of one period: $S_t = \phi(t), t \in [0, T)$. If we build a closed-form approximation, $\hat{S}_t = \hat{\phi}(t), t \in [0, T)$, then we can approximate the state of the system at any time, t' , simply by evaluating $\hat{S}_{t'} = \hat{\phi}(t' \bmod T)$. Provided the error in $\hat{\phi}$ with respect to ϕ is not large, a viewer will not detect the approximation error.

In our tools, approximation of periodic systems involves four steps:

- Determine the period.
- Bound the range of each state variable over one period, to aid in approximating.
- Learn a neural network approximation, $\hat{\phi}(t)$.
- Generate code for the efficient model.

Determining the period of a dynamical system, if one exists, is a common operation in numerical analysis. Fourier analysis may be used, or an approach that searches for return values: places where the function takes on values it has taken before. Our tools take the latter approach.

Given the period, it is possible to determine $\hat{\phi}$, the approximation function, which we represent as a neural network [1]. Neural networks have the advantage of near constant evaluation time for a given network, and are able to approximate well the wide variety of functions our tools may encounter. Specifically, we use a standard feed-forward neural network with two hidden layers and a fixed number of nodes. The network has one input node, corresponding to the time we wish to evaluate at, $t \in [0, T)$, and as many output nodes as there are state variables for the system. Our tools use networks with ten hidden nodes per layer, which is a trade-off between the quality of the approximation and the time taken to evaluate the network. Other network topologies and training schemes could also be used.

The neural network will perform best if the function it is trying to represent has each component in the range $(0, 1)$. In order to re-scale the state variables for a dynamical system to this range, the minimum and maximum possible values for each component must be found. This is done by taking each variable, one at a time, and searching for global minimum and maximum values of that variable over the interval $[0, T)$.

To train the network, N samples of $\phi(t)$ are generated for random times within the period. We then repeatedly apply a standard back-propagation [1] with momentum algorithm on these samples. Each sample is used $N - 1$ times, after which we replace it with a new sample. We do this to reduce the number of dynamical system evaluations, since they are generally slower than a neural network training iteration. The learning process is terminated when the error falls below a user defined threshold, or a maximum number of iterations is reached. A plot of u and \dot{u} as a function of time, and the neural network approximation learned, is shown in figure 2.

The final step is code generation, in which we create a new Java class that evaluates the state of a periodic system efficiently regardless of the time between evaluations. This new system can choose to either use the accurate model or the approximation. It applies the former if the interval between a viewer last seeing the system and the current time is less than 10% of the period, otherwise it applies the latter.

Many approximation strategies other than neural networks are possible. In particular, if the period of the system is short, state variables corresponding to a fixed set of times could be stored, and new state generated simply by interpolation. Such a scheme is directly supported by `Interpolator` nodes in the VRML language. However, in the general case, the neural network approximations we use are smoother and more compact.

3.2 Non-Periodic Systems

As an example of a non-periodic system, we will use the Tilt-A-Whirl (figure 3), an amusement park ride that exhibits highly complex motion despite a simple dynamics description. The ride has seven cars, each attached to a platform on which it is free to rotate. The platforms are driven around a circular hilly track. As the platforms move around the track they tilt so as to remain tangential to




Figure 2: *The exact functions for the parametric position of a roller coaster car on its track, $u(t)$, and its derivative, $\dot{u}(t)$, and the neural network approximation to them, plotted for one complete period (12.13 seconds). The neural network was trained for around 30 minutes to achieve this result. The network approximates u with great accuracy, but does less well with the local maxima and minima of \dot{u} . This is not of concern in this example, because the underlying dynamical system can correct any error based on energy constraints. Currently, our tools use the same network to approximate all the state variables for one system. A separate network could be used for each state variable, which would improve the error at the expense of additional code and evaluation time.*

the surface, which results in complex motions for each car.

Consider a Tilt-A-Whirl that moves out of the view, and then re-enters. If it is hidden for only a short period of time, a viewer can simply extrapolate from the state they saw when it left the view, and hence quite accurately predict the new state. In this case we must use the most accurate model to update the Tilt-A-Whirl's state. However, as the Tilt-A-Whirl is out of view for longer periods, a viewer makes increasingly poorer predictions for what its new state should



Figure 3: A Tilt-A-Whirl amusement park ride.

be. The system can make larger errors in generating the new state without contradicting the viewer. In other words, approximations to the true system may be used, and the approximation error can grow as the time interval out of view grows.

After a Tilt-A-Whirl has been out of view for a long time, a viewer can no longer make use of information from the previous sighting to predict new state. However, a viewer can use their general knowledge of how a Tilt-A-Whirl behaves. To satisfy the viewer's prediction, we must choose a state that is typical for the Tilt-A-Whirl. To represent such typical states, we use a probability distribution over the state space of the Tilt-A-Whirl: states more often seen by a viewer will have a higher probability than states infrequently seen. To generate new state, we simply sample according to the distribution. We will refer to the distribution as the *stationary distribution*³ for the system.

While the preceding discussion is phrased in terms of a Tilt-A-Whirl's behavior, the observations made are typical for any non-periodic system. Various system dependent parameters will change, based on how easy or hard it is to predict the specific system, but the system can be analyzed to find values for these parameters. Our tools do exactly that.

Analysis begins with an accurate model, as for a periodic system, and proceeds through the following steps:

1. **Find the range** of the system: the bounds of its state variables in state space. This allows us to build cell structures over the space and to scale values if required.
2. **Build the stationary distribution** that will be sampled to generate new state when a system has been out of view for a long time.

³The stationary distribution is a term borrowed from the theory of Markov processes, to which our distribution is related [7].

3. **Determine** t_{long} , the time an object must be out of view before we can sample a new state from the distribution.
4. **Build approximations** for generating new state when a system has been out of view for a medium period of time.
5. **Determine** t_{medium} , the time a system must be out of view before we use approximations instead of the accurate model.
6. **Generate code** incorporating the distribution for sampling, the approximations, and control logic for determining which method to use for a given time out of view.

3.2.1 Finding the Range

The range of the system is important because it restricts the region of state space we must concern ourselves with, allowing discrete cell structures to be built on state space. To bound an individual variable, we search forward through time for local minima or maxima for each variable, updating the global minimum and maximum as we go. We stop looking for new local minima and maxima when the global values cease to change significantly. This method is not foolproof — the simulation will not visit regions of state space that are reachable from a different starting point. However, we can be arbitrarily certain of how good the bounds are by tracing a larger number of trajectories from appropriately distributed starting values. We find in practice that small errors in the bounds do not harm the analysis. Also, some variables may be bounded by the user in the input description, particularly angular variables (which lie in $(-\pi, \pi]$).

3.2.2 Building the Stationary Distribution

The stationary distribution is the distribution indicating how much time a long running system spends in any region of the state space. To model the distribution, the reachable regions of state space are divided into constant (user specified) size cells, and a probability, P_i , is attached to each cell i . The result is a discrete distribution on cells, where P_i is the probability that, at a random point in time, the system is within that cell. We assume that the distribution on points within a single cell is uniform.

To build the distribution, we begin with a large number of paths at random and integrate for fixed time-steps, maintaining a counter for each cell measuring how many times a path is in that cell at the end of a time-step. Then,

$$P_i = \frac{count_i}{\sum_i count_i}$$

According to the statistical law of large numbers, the P_i will converge to fixed values as the system is integrated for longer periods of time (assuming a stationary distribution exists). We monitor how much the distribution changes




Figure 4: *The stationary distribution for the Tilt-A-Whirl system. The left image shows a high resolution image of the distribution. Darker points correspond to higher probability, indicating that the system’s state is more likely to take on that value. The right image shows the discrete cell approximation to the distribution. The discrete approximation still captures the overall character of the distribution, but with far smaller data storage requirements.*

between time-steps, and stop when the change becomes small as measured by the L_1 norm.

The discrete cell approximation to the exact stationary distribution performs well in practice, even with quite large cell sizes. Figure 4 shows the stationary distribution for the Tilt-A-Whirl, in which the discrete distribution succeeds in capturing the swirling nature of the exact distribution, but with only a small storage cost.

3.2.3 Determining t_{long}

The sampling threshold, t_{long} , is the period of time that must elapse before a new state may be sampled from the stationary distribution, rather than computed based on some initial conditions. It is equivalent to the time taken for a small region of viewer uncertainty to evolve into the stationary distribution. To see why this is the case, consider what a viewer knows when the object leaves the view. There is some error in this knowledge, which means that the system could be moving on one of several different paths. As time moves on, these paths diverge, until finally the distribution of possible paths looks like any other distribution of paths for the system - the stationary distribution. Because the two distributions are now the same, sampling from one is the same as sampling from the other, and a viewer cannot detect that we sampled from the stationary rather than the exact distribution defined by their knowledge.

To determine t_{long} , we sample a large number of starting values from within a small region of state space, then integrate these paths for fixed time-steps (figure 5). At the end of each step, we check the difference between the distribution




Figure 5: *The convergence of one cell into the stationary distribution, for the Tilt-A-Whirl model. Starting top left and moving clockwise, the plots show the distribution of 5000 paths after 3.07 seconds, 6.15 seconds, 9.23 seconds and 12.31 seconds. The distribution in the lower left is sufficiently close to the stationary distribution to stop testing for this cell. Other cells take up to 24.6 seconds to converge.*

of the paths and the stationary distribution. If these are nearly the same, the total integration time is a candidate for t_{long} . This procedure is then repeated for other starting regions, until enough of the state space has been sampled from. The actual value used is the maximum t_{long} found for any region. Our methods directly examine the propagation of uncertainty, but other approaches based on more theoretical considerations [5] may also be applicable.

3.2.4 Building Approximations

The approximation functions we build in this step will be used to generate new state quickly, with some error allowed. After some short period of time, they must be cheaper to evaluate than the most accurate routine supplied by the user, and we want the cost of evaluating them to grow more slowly than the time period over which they are evaluating. For this task we use neural networks, similar to those used for approximating periodic functions.

In this step we generate several neural networks, each of which evaluates over its own fixed time interval, Δt_i . As input, each network takes the state of the system at time t , so there are as many input nodes as state variables. On output, each produces the state at time $t + \Delta t_i$. One network evaluates a function over a period of half the sampling threshold, $t_{long}/2$. The next function evaluates over half this time, the next over half of that and so on, stopping at a network that evaluates over either a user defined minimum step or the period

of any forcing functions, whichever is greater. By chaining neural networks together, we can evaluate to within a small distance of any time interval, which we will reach exactly using the accurate model.

Our structure of networks has the following advantages:

- We expect the cost of evaluating networks to grow at a slower rate than the time interval for which they evaluate, resulting in computational savings over the cost of evaluating one network over many steps.
- We can build the network's dependence on time into the network itself, rather than making it an input parameter. This significantly simplifies the network, and allows lower errors for the same size net.
- We can tolerate larger errors in networks that evaluate over longer times, without sacrificing accuracy in networks that will evaluate over short periods.
- We can learn networks concurrently, with significant improvements in training time and efficiency.

The samples the neural networks are trained on are distributed according to the stationary system (due to the method we use to generate them). This will result in the networks having lower error rates in regions of high probability, and higher rates in regions of lower probability (figure 6). This is acceptable, because the error will tend to be inversely proportional to the likelihood of a viewer seeing the error. We terminate learning if the error falls below a user defined threshold value. We grow the network by adding five new nodes per layer each time its learning rate slows. We also force termination after a fixed number of cycles if the network has not reduced its error to an acceptable level.

We could use other approximating functions, such as radial basis functions, wavelets or splines, which have the advantage of elegant subdivision schemes, but lack the generality and ease of fitting offered by neural networks. For optimum approximation, the best method would be chosen based on how each performed on the target function.

3.2.5 Determining t_{medium}

To determine how long before the approximation may be used instead of the true evaluation routine, we may simply find the point at which it becomes more efficient to approximate, provided the approximation error is sufficiently low. The neural network learning procedure ensures that the error in the approximation function for the shortest evaluation time network is within a viewer's ability to predict, and we assume that neural networks are always cheaper to evaluate than the accurate system, so we simply set t_{medium} to the smallest evaluation time of the networks we have learned.




Figure 6: *An example of the Tilt-A-Whirl equations of motion approximated by a neural network. Intensity indicates function value: lighter is higher valued. The network shown is attempting to learn the change in orientation and change in velocity of a Tilt-A-Whirl car over a 6.15 second interval. The left column plots the change in position as a function of initial conditions, and the right column plots the change in velocity. The top row shows the true function, the middle images show the function learned by the network, and the bottom row shows the difference image, masked by the stationary distribution shown in figure 4. In each frame, the initial orientation increases from $-\pi$ to π along the horizontal axis, and the velocity increases vertically from -4.21 to 3.59 . Note that we are not concerned with errors that are masked out by the stationary distribution, because these errors will never be seen by a viewer. Such regions include the top left and bottom left corners of each frame.*

3.2.6 Code Generation

The code generated for non-periodic functions allows efficient evaluation of new dynamic state over any time interval and within a viewer's ability to detect errors. The components of the new model are:

- A representation of the stationary distribution and code to sample from it.
- Code to evaluate the various neural networks, and a wrapper function that determines the set of evaluations required to step forward a given amount in time.
- Control logic that examines the difference between the desired evaluation time and the last time the system was seen by the viewer. If the difference is greater than t_{long} it samples new values. Otherwise, if the difference is greater than t_{medium} , it uses a neural network approximation. Otherwise, it uses the accurate model to generate state.

4 The Runtime Layer

The tools for generating approximations for culling produce a description of the dynamical system that can efficiently evaluate the state at any given time regardless of the period between evaluations. The aim of the runtime layer is to provide a standard interface between such systems, written in Java, and a VRML world. From the VRML side, a dynamical system is an **EXTERNPROTO** object, reducing to a minimum the complexity of incorporating a model into a larger world. From the Java side, it is only necessary to supply the dynamics evaluation routine in a standard format.

A runtime environment is generated for each dynamic object, consisting of a VRML file and a Java class file. Together, they must perform the following tasks:

- Store the geometry of the animated objects.
- Link the state variables of a dynamical system to transformations of the objects.
- When the VRML browser indicates a new frame⁴, the runtime system must obtain dynamic state values for that frame.
- Track the visibility of the system, and turn off evaluation of dynamic state (cull this system) if not in view.

⁴This is achieved through `TimeSensor` events, which we assume a browser sends at least once per frame

The geometry of the animated object is stored within the VRML file, in standard VRML format. The VRML file also contains a **Script** node to interface to the Java, and it defines all the necessary names and **ROUTES** for linking the output of the script to transformations of the geometry.

To manage dynamic state, the runtime environment maintains buffers of state variables evaluated at fixed time intervals. Intermediate values required for rendering are interpolated between buffered states. We use linear interpolation, although higher order schemes are possible. Buffering state is advantageous because it provides constant frame rate interface to the renderer, while allowing the underlying dynamic system to compute values at any rate and time-step. It also makes it possible to have useful values ready in the buffer if the object leaves the view, and then re-enters very soon after.

When the runtime system receives a request for new state, there are two possibilities, depending on whether the value is already buffered. If the value requested does not appear in the buffer, the dynamical system is evaluated twice, once for each of the values bracketing the requested time. If the values are already in the buffer, the system is still evaluated in order to fill future slots in the buffer. In practice, the system is repeatedly evaluated to fill the buffer until it signals that enough work has been done for one frame.

To track visibility, we use **VisibilitySensors** within the VRML file. These sensors send events each time a bounding volume enters or leaves the user's view. In turn, these events cause the Java class to mark individual dynamical systems that control the object as visible or invisible. Invisible objects will not be evaluated for new state. The VRML specification does not allow a user to explicitly activate or deactivate Java scripts, so we use an automatically generated script to ensure that events are not sent to culled scripts.

Each runtime environment is unique to its model. To create the runtime, the user provides a file describing the geometry to be animated and associating dynamic variables with transformations. A program then completes a template to create the runtime files. The file format used is simple enough to generate by hand, but the easiest method is through a modeler, which we describe next.

5 Attaching Dynamics to Geometry

The rigid-body modeler allows a user to load objects described by VRML's **PROTO** mechanism, display them, and build transformation hierarchies consisting of objects, rotations, translations and bounding boxes. The values used for the transformations may be animated, thus creating dynamic models. As output, the modeling program produces the runtime layer as described above, or it can save an intermediate file format. The interface is shown in figure 7.

In describing systems for the modeling process, we distinguish between output variables and state variables. State variables are the set of values required to describe the system completely at any time, whereas output variables correspond directly to geometric transformations. The output variables are derived from the state variables through a user defined function. For example, the roller

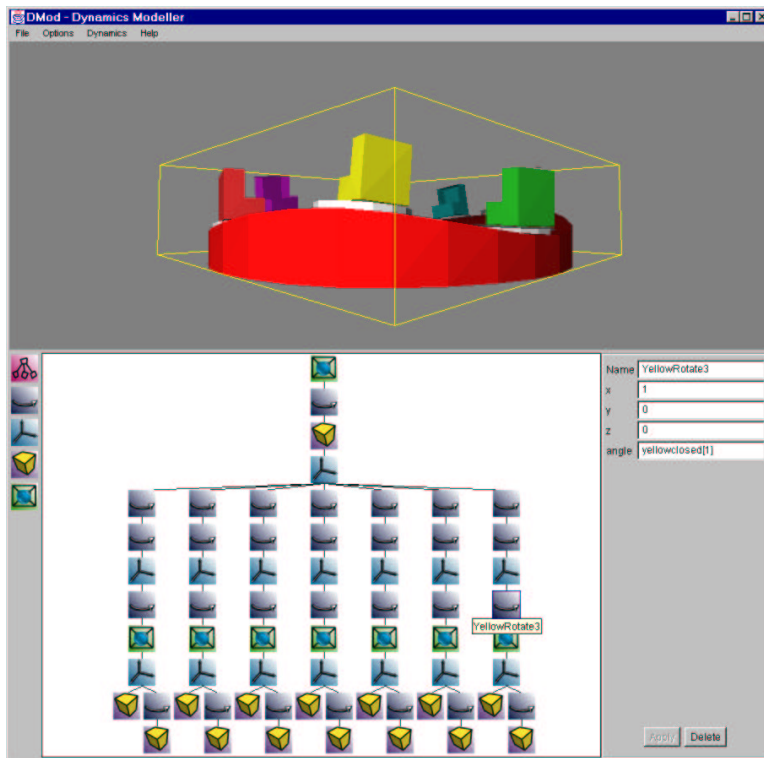


Figure 7: *The interface for our dynamic transformation modeler consists of three regions. In the top panel the current geometric arrangement of the system is displayed, including its bounding volume. The transformation hierarchy for the system is displayed in the lower left panel. A simple drag-and-drop interface is used to edit the hierarchy. The available nodes, shown on the far left are (top to bottom): group, rotation, translation, VRML object and bounding volume. By clicking on a node in the tree, a user can edit the fields of that node in the lower right panel, entering a constant, or a variable to animate the field. The user is shown here editing a rotation node, specifying a variable as the angle through which to rotate.*

coaster model has two state variables: the parametric position on the track and the parametric velocity along the track. The output variables are the actual position in world space, and the various rotations to align the car with the track, derived by evaluating the track spline at the position indicated by the state variables.

Each field of geometric transformation is specified either as a constant value, or as an element of a field variable. Each field variable is an array corresponding to the output variables for a dynamical system, which the user specifies as a Java class name when defining the field variable. The class for field variables

must implement an interface that defines functions for evaluating the state at a given time, and for setting the values of output variables. The approximation tools described above produce such classes, but users are also free to define them directly — the format is simple and compact. Hand authoring is desirable for systems expressed as closed-form functions of time, which do not benefit from the approximations described above yet are still important for modeling dynamics.

Within the modeler, a user can preview the effect of the dynamical system on the model. Such a preview takes place within a Java environment, as opposed to the VRML browser environment, which makes debugging significantly easier.

The bounding boxes in the hierarchy are used to tell the system which dynamics may be culled. Each bound has associated with it a set of variables whose visual effect is contained within the bound. If the bound is not visible the dynamics for those variables will be culled. The modeler can automatically determine this bound, by running the system within the modeler and examining the maximum extents of geometry over time.

We emphasize again that the dynamics are largely independent of the geometry they are attached to. There is some dependence if the simulation is to be physically plausible, such as lengths of geometric objects appearing as parameters in the dynamical system, but as parameters they are readily made available to a user, and don't change the structure of the underlying equations. This allows re-use of parameterized dynamics with different geometries, and vice-versa. More importantly, it makes possible a library of dynamical systems, each with efficient cullable code, which could be used by authors in the same way 3D geometry libraries are used today.

6 Speedup Results

To measure the performance improvement while culling dynamics, we studied the time spent computing the dynamics for each rendered frame of a simulation. The models used in our tests are shown in figure 8. We began with one instance of each in the world, then added additional objects up to a maximum of 35 (five of each example). In each test, the average frame time was measured for a viewpoint animated such that the center of view moved in a circle around the world while the view direction oscillated through a 90° angle. Rides were added so that the density of rides in the world was approximately constant. With culling turned on, the dynamics for a ride were computed only if the ride was visible, and we used the models generated by our software to ensure fast, consistent evaluation. With culling off, the dynamics for all the models were computed for every frame using the accurate model for the system. The geometric rendering was not affected by the culling — we assume the browser was culling geometry against the view volume.

Table 1 presents the timing values recorded in our experiments. Our results demonstrate that the average time per frame is roughly linear with respect to the average number of systems in view, because we only perform computation




Figure 8: *Fairground rides modelled with our tools. Shown are: The octopus, front left; The pirate ship, front right; Two version of a double pendulum, center front and back; The divebomber, back right; and the roller coaster and tilt-a-whirl.*

Total	Cull On	Cull Off	# In View	% in View
7	0.048s	0.126s	1.42	20%
14	0.090s	0.249s	2.94	21%
21	0.142s	0.394s	4.75	23%
28	0.208	0.550s	7.27	26%
35	0.280	0.727	9.12	26%

Table 1: *Average time per frame with and without culling, and the average number of models in view, for increasing numbers of models. The time per frame with culling on grows approximately linearly with the number of models in view, and the frame-rate speedup is roughly constant (figure 9), as expected with our set up.*

for objects in view, and the browser performs geometric culling against the view volume. Figure 9 plots the speedups obtained by culling over a world that does no culling. The speedups obtained are around 2.6, and remain roughly constant as the number of systems increases. We expect this result, as the percentage of systems in view remains approximately constant. We did not achieve the $4\times$ speedup one might expect (only around 1/4 of the systems were in view at any time). This is due in part to the fact that our underlying system must perform some computation to check whether each system is in view, as well as the overhead of Java scripts.

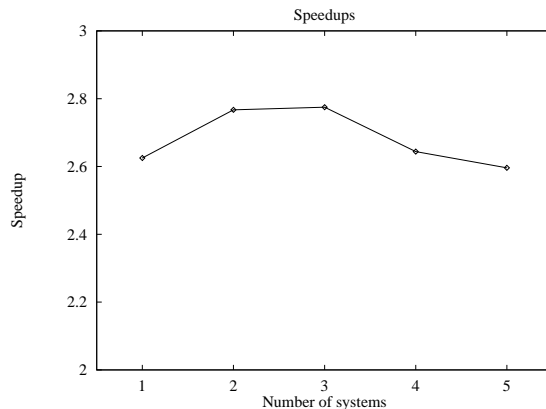


Figure 9: *The speedups obtained plotted as a function of the total number of models in the world. Note that the vertical axis does not start at zero.*

7 Future Work

One significant extension to our approximation software is to add the capacity to handle state machines and hierarchies of systems. The optimization approach is similar, and the range of systems that can be modeled would be greatly increased. For example, the Tilt-A-Whirl ride would be able to stop to let off and collect passengers, rather than run indefinitely.

For novice users, an authoring tool would ideally hide any equations of motion for the system from the author. Such a tool would approach modeling from the point of view of geometric constraints and common forces (such as gravity or motors). The modeler would then infer the dynamics and generate the accurate model required by our current system. With this architecture the optimization process is carried out as a subsystem of the modeler, and the dynamics need never be explicitly stated by the user.

Some aspects of our work could be simplified by incorporating appropriate semantics into the VRML specification. In particular, our runtime layer uses special scripts to ensure that timer events are only sent to systems in view. The VRML browser could be supplied with all the information necessary to do that itself: scripts require bounding volumes with the semantics that scripts need not receive events when their bound is invisible. In many cases this information could be inferred by examining the bounding volumes for geometry that is influenced by events from the script.

Of interest to the VRML community is the interaction of culling with multi-user networked environments. A conservative culling approach for multi-user worlds would keep track of the last time each system was seen by any viewer, and apply the same tests used in this article to determine how the system should be evolved when it is next seen. More aggressive approaches might allow different viewers to see different things, based on the viewers' ability to communicate. In particular, as long as each viewer perceives the same events, their accounts will

agree. If the actual dynamic motion that led to the percept is slightly different, that is unlikely to lead to noticeable inconsistencies.

8 Conclusion

We have presented a set of tools for authoring efficient dynamic models in VRML and Java. They make use of novel techniques for automatically generating dynamical models that may be culled when not in view. Our interactive modeler allows geometry and dynamics to be combined by associating geometric transforms with dynamic state variables.

Our results show that large numbers of models can be included in a VRML world without sacrificing frame rate. We achieve this by performing only work that is necessary to ensure a consistent environment for the viewer. Our tools make it possible for those inexperienced with dynamics to achieve similar results, while strongly encouraging the development of a library of dynamics for use with varying geometry.

9 Acknowledgments

Our work was performed on machines donated by Intel, and funded by the Office of Naval Research grant ONR-MURI-N00014-96-1200. We used the Worldview for Netscape Navigator browser running under Windows NT on a 200MHz Pentium Pro. We used Sced, a public domain constraint-based modeler, for the geometric modeling.

References

- [1] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.
- [2] Stephen Chenney and David Forsyth. View-dependent culling of dynamic systems in virtual environments. In *Proceedings 1997 Symposium on Interactive 3D Graphics*, pages 55–58, April 1997. Providence, RI, April 27-30.
- [3] Stephen Chenney, Jeffrey Ichnowski, and David Forsyth. Efficient dynamics modeling for VRML and Java. In *VRML 98, Proceedings 1998 Symposium on the Virtual Reality Modeling Language*, 1998. Monterey, CA, February 16-19, To Appear.
- [4] Radek Grzeszczuk, Demetri Terzopoulos, and Geoffrey Hinton. Neuroanimator: Fast neural network emulation and control of physics-based models. In *Computer Graphics, Proceedings of SIGGRAPH 98*. ACM SIGGRAPH, 1998. To Appear.

- [5] John Guckenheimer and Philip Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, New York, 1983.
- [6] Jessica K. Hodgins and Deborah A. Carlson. Simulation levels of detail for real-time animation. In *Graphics Interface '97*, pages 1–8, 1997. Kelowna, BC, Canada, 21-23 May 1997.
- [7] Chieh Su Hsu. *Cell-to-cell mapping: a method of global analysis for nonlinear systems*. Springer-Verlag, New York, 1987.
- [8] M. N. Setas, M. R. Gomes, and J. M. Rebordão. Dynamic simulation of natural environments in virtual reality. In *SIVE95: The First Workshop on Simulation and Interaction in Virtual Environments*, July 1995. University of Iowa, Iowa City, IA.