

Specification

For some applications of VRML (such as database visualization, scientific visualization, and possibly shared, multi-user, distributed worlds), it is desirable to allow the creation and modification of VRML scenes. We believe this need with a proposal that allows the VRML developer to either take complete control over changes to the scene, which may limit the opportunities a browser has for optimization, or to either limit their control to a portion of the scene or to limit their changes to valuable objects. This change communication model is combined with a new prototyping capability that allows the encapsulation and re-use of objects, behaviors, or both.

We are proposing a model for describing how changes are communicated between objects in the VRML scene. VRML's existing notion of a field of an object is extended to allow imports and exports; imports and exports can be thought of either as communication channels or messages that can be sent or received.

3. Other input sources, accessed through some scripting logic language, probably running asynchronously to the browser.
2. Time.
1. User input from some device, translated into the 3D world by the browser.

The essence of simple behaviors is changes to the world over time. With a single spectator interacting with the world, we identify three sources of change:

Overview

We will not know if our goals and constraints have been met until we have implemented this system in both a browser and an authoring system. However, we have considerable experience with both browsers (WebSpace Navigator) and authoring systems (WebSpace Author) for VRML, and believe that this design will strike a good balance between speed for the VRML speciator, power for the VRML hacker and ease of use for the VRML artist.

We expect VRML to evolve into a multi-user shared experience in the near future. Allowing VRML browsers might eventually need to support synchronization of changes to the world, complete collaboration and communication in interactive 3D worlds. We have attempted to anticipate the needs of multi-user VRML, in our design, considering the possibility that VRML browsers might eventually need to support synchronization of changes to the world, locking, persistent distributed worlds, event rollback and dead reckoning in the future. We will not know if our goals and constraints have been met until we have implemented this system with both a browser and an authoring system. However, we have considerable experience with both browsers (WebSpace Navigator) and authoring systems (WebSpace Author) for VRML, and believe that this design will strike a good balance between speed for the VRML speciator, power for the VRML hacker and ease of use for the VRML artist.

We must allow programmers to seamlessly extend VRML's functionality, by allowing them to create arbitrary scripts/applets/code that can then be easily re-used by the non-programmer. Multi-user portability. We expect VRML to evolve into a multi-user shared experience in the near future. Allowing VRML to perform most of the tasks necessary to create an interesting, interactive VRML world using graphical user interface. We believe that VRML will not be successful until we assume that sophisticated VRML authoring tools will be created, and wish to make it possible for people that are not interested in programming to create VRML content.

Power. We must allow programmers to seamlessly extend VRML's functionality, by allowing them to create arbitrary scripts/applets/code that can then be easily re-used by the non-programmer. Multi-user portability. We expect VRML to evolve into a multi-user shared experience in the near future. Allowing VRML to perform most of the tasks necessary to create an interesting, interactive VRML world using graphical user interface. We believe that VRML will not be successful until we assume that sophisticated VRML authoring tools will be created, and wish to make it possible for people that are not interested in programming to create VRML content.

Scalability. We believe that speed is a key to a good interactive experience, and that it is important to design the system so that VRML browsers will be able to optimize the VRML world. Our goal is to allow the creation of very large virtual worlds. Any feature which limits scalability is unacceptable. And in several areas of our design we have purposefully made certain things difficult to encourage the creation of scalable VRML worlds.

Performance. We believe that speed is a key to a good interactive experience, and that it is important to design the system so that VRML browsers will be able to optimize the VRML world. Our goal is to allow the creation of very large virtual worlds. Any feature which limits scalability is unacceptable. And in several areas of our design we have purposefully made certain things difficult to encourage the creation of scalable VRML worlds.

Goals. Our design was guided by the following constraints (in rough order of priority):

- Provide the rich, interactive experience that will enable a complete new set of applications of VRML.
- Behaviors, the ability to create animated 3D objects, and the addition of sound into 3D worlds will provide the user with the ability to interact with intelligent objects with simple behaviors.
- Possibility in VRML 1.0. We believe that the user a much richer experience than is possible in VRML 1.0.

Section Graphs

Rob Myers
Chris Martin
Chee Yu
Gavin Bell, gavint@sgi.com

Simple behaviors for VRML

The node names specified in the input and output declarations must be DEFed inside the prototype implementation. The first node DEFed in lexical (not traversal) order will be exported. It is an error (and results are undefined) if there is no node with the given name, or the first node found does not contain a field of the appropriate type with the given field name.

Specifying the type of each input or output in the prototype is intended to prevent errors when the scene graph rooted by node is exported inputs and outputs inside the scene graph given by node. The input and output declarations export inputs and outputs with the same name to be exported with unique names. Prototypes do not have fields, and fields inside the implementation may implement a name for each input or output to allow several inputs or outputs with the same name to be specified instead of prototypess.

A prototype is NOT a node; it merely defines a prototype (named *typeName*) that can be used later in the same file as if it were a built-in node. The implementation of the prototype is contained in the scene graph rooted by node. The implementation may not be DEFed or USEd.

```
proto typeName [ input fieldType name name ]
    is nodeName.fieldName nodeName.fieldName ...
    output fieldType name name ...
node ( ... )
```

A prototype is defined using the **PROTO** keyword, as follows:

PROTO

Prototyping is a mechanism that allows the set of node types to be extended from within a VRML file.

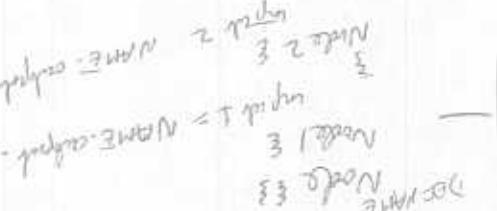
Prototyping

If an input is not connected, it may be given a constant value. Inputs that are connected may not be given a value. Inputs that are not connected and are not given any value will have a default value.

Loops are well-defined, deterministic, and useful in this proposal.

```
DEF NAME Node [ input useName, output useName ]
    output = USE NAME, output
DEF NAME Node { }
```

DEF/USE may also be used to create loops:



An output may be connected to several inputs (*fan-out* is allowed), simply by using VRML's **exists** mechanism:

Note: We considered allowing fan-in, and defining automatic rules for combining multiple outputs into one value. Note that a "last-one-wins" rule does NOT work, because "last" is ill-defined; outputs may change at exactly the same time. We believe it better to require authors or authorities to handle situations in which fan-in is desirable by defining logic nodes with multiple inputs that handle situations in which fan-in is undesirable by defining logic nodes with multiple outputs that combine the multiple input streams.

An input may have only one thing connected to it (fan-in is not allowed); the syntax chosen enforces this rule.

```
Node { input = Mode( ), output }
```

An input of some node is connected to an output of some other node using the following syntax:

We propose that many of the current VRML nodes be redefined such that their fields are declared to be inputs. An input is merely a field which may be connected to an output. Connecting an input to an output establishes a communication pathway between two objects; the object containing the input will be told whenever the output changes, and may then change its own outputs, render itself differently, etc. Inputs and outputs are written using the same format as VRML fields.

The following describe in detail what is proposed:

Connections

Prototype declarations have file scope, and prototype names must be unique in any given file.

A prototype is instantiated as if it were a built-in node. A prototype instance may be DEFed or USEd. For example, a simple chart with variable colors for the leg and set might be prototyped as:

```

PROTO TwoColorChart { input SPColour LegColour IS Leg, diffuseColor,
    input SPColour SeatColour IS Seat, diffuseColor }
    Separator {
        Cube { ... }
        DEF seat Material { diffuseColor .6 .6 .1 }
        DEF leg Material { diffuseColor .8 .4 .7 }
    }
    Separator {
        Cube { ... }
        DEF seat Material { diffuseColor 1 0 0 }
        DEF edgeMaterial { LegColour 1 0 0 }
    }
    USE edgeMaterial # regular DEF/USE rules apply
}

```

Note: **PROTO** sort of gives people their non-instantiating DEF: **PROTO foo () Cube ()** is roughly equivalent to DEF **Cube ()**, except that **foo** is now a type name instead of an instance name (and you say **foo ()** to get another cube instead of USE **foo()**). Smart implementations will automatically share the unchanged stuff in prototype implementations, so the end result will be the same.

What if we wanted a prototype that could be instantiated with arbitrary geometry? For example, we might want to define a prototype chart that allowed the geometry for the legs to be defined, with the default (perhaps) being a simple cylinder.

VML 1.1 will include the SFNode field **type**—a field that contains a pointer to a node. Using

SFNode, it is easy to write the first part of the PROTO definition:

PROTO Chair { input SFNode # input and inserting it into the scene. This can be accomplished with a new node,

... but then we get stuck when we try to define the IS part of the prototype. We need some way of taking an SFNode input and inserting it into the scene. This can be accomplished with a new node,

the NodeReference node:

NodeReference { nodeToUse Sphere { } }

PROTO Chair { input SFNode # input SFNode (NULL is valid syntax for SFNode)

nodeToUse NULL # input SFNode (NULL is valid syntax for SFNode)

NodeReference { nodeToUse Sphere { } }

Functionally, NodeReference is a "do-nothing" node—it just behaves exactly like whatever

nodeToUse (unless nodeToUse is NULL, of course, in which case NodeReference does nothing). For example, this would be a verbose way to add a Sphere to the scene:

NodeReference { nodeToUse Sphere { } }

For example, a Logic node which is a latch that toggles its output value whenever both of its inputs are TRUE can be written like:

The language of the script consists of fields used in square brackets and separated by commas. The script field contains either the location of the script or the actual script itself. The language, field specifies modifiers, types and names of fields used in the keyword fields followed by a list of the access fields, a script and the script language. The field description is written just after the opening tag.

The Logic node encapsulates a pure executable function in VML. It consists of a description of its

Logic: scripting

We need to think about scalability when using nested EXTERNPROTOs. EXTERNPROTOs don't have bounding boxes specified like WWWlinks, and they might need them. I'm starting to think that we might need to add bounding boxes instead of having them only on

WWlinks; with annotations possible, pre-specifying maximum-possible bounding boxes could save a lot of work recalculation bounding boxes as things move.

Note: The rules about allowing exporting only from files that contain a single PROTO declaration are consistent with the WWWlinks rules; until we have VML-aware protocols that can send just one object or prototype declaration across the wire. I don't think we should encourage people to put multiple objects or declarations in a single file.

In this case, the implementation of the prototype is found in the given URL. The file pointed to by that URL must contain ONLY a single prototype implementation (using PROTO). That prototype is then given the name type name in this file's scope (allowing possible naming clashes to be avoided), it is an error if the name type name is passed in the EXTERNPROTO is not a subset of the input output

that rules about allowing exporting only from files that contain a single PROTO declaration are consistent with the WWWlinks rules; until we have VML-aware protocols that can send just one declaration specific in URL.

externproto type name [input fieldtype name] output fieldtype name . . .

URL

A second form of the prototype allows prototypes to be defined in external files:

EXTERNPROTO

Something else to think about: should a prototype be allowed to expose the outputs or use the inputs of the contents of an SNode that is passed in? E.g. could I define a prototype that took a node as input, assumed that the node had a "foo" output, and published the output of that node as an output of the prototype?

A really smart browser will figure out that nobody is using the optimized scene and may free it from memory.

Browsers that use optimized scene graphs can implement NodePreference such that whenever nodeToUse changes an optimized scene is created. When rendering, the optimized scene will be used instead of the unoptimized scene.

Browsers that want to maintain a different internal representation for the scene graph can implement NodePreference so that nodeToUse is read and the different internal representation is generated.

Optimizations might also be performed at the same time.

The NodePreference node has nice, clean semantics, and allows a lot of flexibility and power for defining prototypes. It also has some nice implementation side effects.

Note that SNode fields follow the regular DEF/USE rules, and that SNode fields contain a pointer to a node; using DEF/USE an SNode field may contain a pointer to a node that is also a child of some node in the scene graph, that is pointed to by some other SNode field, etc.

Chart (LegGeometry USE LEG)

It might also make sense to share the same geometry between several prototype instances; for example, you might do:

Chart (LegGeometry Separator (Coordinate3/IndexFaceSet/etc)

other field type. Using the Chart prototype would look like:

more units (such as meters) shows a certain value to be given to the input of a prototype, just like any

This allows the creation of:

```
language "C"
{
    state = state;
    value = state;
}
script .c (spawned is
    state = state;
    SPPool state;
    output SPPool value;
    input SPPool t
}

isdoneon FALSE isdownon FALSE value FALSE state FALSE
```

The logic node is scripting-language neutral, and this document describes only the properties required of any scripting-language, and describes in general terms what kind of functionality the script language provides.

Logic nodes can maintain either state that the browser knows about ("public state", which is implemented simply as fields of the logic node) or internal state that the browser is not aware of. Scalable implementations will load and unload logic nodes as different parts of the virtual world are loaded and unloaded into and out of memory (we assume that virtual worlds will be larger than available disk or memory). To maintain the illusion of a continuous virtual world, when unloaded from the world browser must remember that the lights in that room are on even if the room is unloaded again. For example, if the user walks into a room, turns the lights on, then walks out of the room, the browser must remember that the lights in that room are off even if the room is unloaded again. To make this saving/restoring fast and small, authors should try to minimize the amount of public state. Any state that can be easily re-created should be stored as private state inside the logic node. For example, a logic node that outputs the last 100 trade prices for a specific stock on the stock market might have public state of:

```
Logic {
    tickersymbol "SGI";
    stockserver "stocks://www.cgi.com/stockfeed";
    ... outputs an array that is last 100 prices
```

Given the stock server and stock symbol, the output can be re-created, so this is all the public state needed.

Evaluation Order

To guarantee that a behavior produces similar results on different VRML implementations, rules are defined for determining the evaluation order of logic nodes.

Authors need a well-defined, consistent mental model of how behaviors interact. The challenge is to describe this model precisely enough so that authors can be confident that they are creating behaviors that will work across implementations, but also to allow enough freedom so that different implementations of behaviors are possible.

Implementations will present the results of behaviors by sampling the behavior's output at a particular point in time and rendering the result. This sampling of continuous time is an implementation detail; the author has no control over it (implementations may choose to sample behaviors just before or during rendering, may choose to run a separate simulation thread that samples behaviors more frequently than the render time, or make choices to turn off behaviors to save bandwidth). The mental model we have chosen consists of continuous changes over time and discrete events that occur at particular points in time. "Time" refers to an idealized notion of time corresponding to the everyday notion of time ("wall-clock" time in other systems).

Implementation details of the results of behaviors by sampling the behavior's output at a particular point in time and rendering the result. This sampling of continuous time is an implementation detail; the author has no control over it (implementations may choose to sample behaviors just before or during rendering, may choose to run a separate simulation thread that samples behaviors more frequently than the render time, or make choices to turn off behaviors to save bandwidth).

Editing the scene

Exactly what this looks like in the scripting language depends on which scripting language is being used. And exactly what this API looks like in the implementation of the browser and controls depends on which language it depends on.

Compare this with the synchronous inputChange method described above:

```
do forever:
    wate for input from somewhere in the real world
    sync with browser
    get state of inputs/fields
    perform some calculation
    set state of outputs/fields
    sync with browser
    do forever:
```

If running asynchronously, the evaluation loop will probably look something like:

```
start.
anything allocated in the start method, including terminating any asynchronous processes
will probably have an inputChange method that communicates with the asynchronous
process, telling it that something has changed. If an asynchronous logic node is not being
used, then the inputChange method will typically look like:
```

```
set state of outputs/fields
get state of inputs/fields
perform some calculation
set state of outputs/fields
end
```

When one or more inputs of a logic node change, some kind of notification/evaluation method must be called. The inputChange method is synchronized, and is expected to perform some calculations, set outputs, and return. An asynchronous logic node with inputs will probably have an inputChange method that communicates with the asynchronous process, telling it that something has changed. If an asynchronous logic node is not being used, then the inputChange method will typically look like:

```
logic node care about both changes or only the final value?
second, if the input of a logic node changes twice during that I/O of a second, does the
implementation runs behaviors just before rendering, but rendering takes I/O of a
final values of its inputs or whether it cares about immediate changes (for example, if
the browser should expect spontaneous changes and whether the script cares only about
that information back to the browser about whether or
method will need some way of returning information to the script or
some start or init method, to allow the script to startup an asynchronous process. The start
script the browser/script API will need at least the following methods:
```

- gets the value of an input
- sets the value of a field
- sets the value of an output

The evaluation model remains the same. The asynchronous process must be synchronized with the browser whenever it does any of the following:

To support sources of input from outside the virtual world (logic nodes that spontaneously produce logic nodes function very much like the built-in Sensor classes). It may be desirable to allow logic nodes to run as asynchronous processes. In this case, the output of the browser/script API should probably allow a script to specify "Please evaluate me whenever this specific input changes"; for example, if a script wants to know when a button goes up a script can be lazy about the last mouse click. We will realize significant savings if we can be lazy about evaluating scripts.

Asynchronous Logic

Scripts with internal state cannot be optimized as much as scripts whose outputs depend solely on their inputs. The browser/script API should probably allow a script to be lazy about evaluating scripts unless they are modified to do that without severely restricting timing implementations?

It is possible for values to change while the same thing... Loop semantics also need to be clearly stated; is and compare values to accomplish the same thing... Loop semantics also need to be clearly stated; is their scripts - to know if something does/dosen't need to be recomputed. But they could just store change at the same time, is that a useful thing to know? It could potentially allow authors to optimize be allowed to ask which inputs changed to cause the evaluation? Since, potentially, all inputs might write logic nodes. For example, when one or more inputs of a logic node change, should an author

However, doing so will result in non-predictable behaviors, and should be discouraged. It is possible for authors to create behaviors that are dependent on how often behaviors are sampled.

the fields center and size (an object-space axis-aligned box).

The VolumeProximitySensor node reports when the camera enters and leaves the volume defined by

VolumeProximitySensor

write a Logic node with internal state that figures out...).
might also be useful, but I'm not convinced they're useful enough to add (and besides, you could
wherever it is in the world). Position and orientation at the time the user entered/exited the world
(authors can create proximity sensors that enclose the entire world if they want to track the viewpoint
is outside the region would kill scalability and compatibility).
Issue: there are some other things we might like to know - current position and orientation while the
viewpoint is inside the region could be very useful. Providing position and orientation when the user
is outside can create proximity sensors that enclose the entire world if they want to track the viewpoint
types of proximity sensors: VolumeProximitySensor and PointProximitySensor.

While the camera remains within the region, the output is in remains TRUE. There are two
volume. The camera crossed this volume, while the exit output remains the last time the camera exited this
sensor is triggered by something in the input to TRUE. The enter output contains the last time that
proximity sensor can be made interactive at any time by enabling its enabled input to FALSE. An inactive
proximity sensors are nodes that detect when the camera enters and exits a defined region. A

Proximity Sensors

Sensors: Input

Exit asynchronous process cleanly.
end

Tell the browser we're done making changes
Make changes to the scene

would be bad)

Synchronize with the browser (making changes while it was in the middle of rendering
Wait until the server sends information about changes to the scene, then:

Asynchronous process

that the asynchronous process can:
the server and set the geometry output to that scene. Also maintain a pointer to the scene so

field or input of the Logic node), construct an initial scene graph based on information from
More implementation notes: for this to work, NodeReference will have to notice that screen might be a

start

And the Logic node would have methods that did:

```
nodeReference = Logic ( ... ) . geometry  
nodeReference ( .  
Separator ( .
```

Look like:

For example, a Logic node might receive scene graph changes from a server. The VRML file would

NodeReference nodes point to are constantly changing.
since browsers with different internal representations will take a significant hit if the scene's that
authors should still try to minimize which parts of the scene they don't allow to be changed arbitrary.

NodeReference is changing, and not bother doing the optimization if it is changing too often.
the next time it is rendered. Really smart optimising browsers might keep track of how often a
nodeToUse points to has been changed, and may need to re-optimize (or regenerate) the internal rep-

More implementation notes: for this to work, NodeReference will have to notice that screen might have

arbitrary changes to one or more parts of the scene graph.

Because an SFNode field stores a pointer to a node, a Logic node may decide to maintain a pointer to
a node it is outputting as part of either its public or internal state. This allows a Logic node to make

A Torus prototyped this way will look exactly as if it is a built-in node. New property nodes (based
on existing properties) can also be defined in this way.

```
.geometry  
writes Separator pointer to geometry output  
Separator/Coordinate3/IndexedFaceSet,  
script "reads radius/zRadius, generates  
radius 1 radius 1 geometry NULL # Default values  
fields { input SFRNode geometry }  
fields { input SFRNode radius, input SFRfloat radius2,  
nodeReference = DEF GeneratorLogic {  
    input SFRfloat radius1 IS generatorRadius, radius2  
    nodeReference {  
        input SFRfloat radius1 IS generatorRadius, radius2 IS generatorRadius, radius2  
    }
```

interesting things possible; for example, to define a Torus node you might
can pass in the active region, containing this with the proxy URL mechanism makes some pretty

```

FILE FORMAT/DEFAULTS
    ClickSensor {
        input enabled TRUE # SBOOL
        output center 0 # SFTIME
        output exit 0 # SFTIME
        output enter 0 # SFTIME
        output isOver FALSE # SBOOL
        output release FALSE # SBOOL
        output scroll 0 # SFTIME
        output touch FALSE # SBOOL
        output volume 0 # SFTIME
        output x 0 # SFTIME
        output y 0 # SFTIME
        output z 0 # SFTIME
        radius 0 # SFLOAT
        size 0 0 0 # SFVEC3F
        center 0 0 0 # SFVEC3F
        output isOver 0 0 0 # SFVEC3F
        output scroll 0 0 0 # SFVEC3F
        output touch 0 0 0 # SFVEC3F
        output volume 0 0 0 # SFVEC3F
        output x 0 0 0 # SFVEC3F
        output y 0 0 0 # SFVEC3F
        output z 0 0 0 # SFVEC3F
        output volume 0 0 0 # SFVEC3F
        output scroll 0 0 0 # SFVEC3F
        output touch 0 0 0 # SFVEC3F
        output x 0 0 0 # SFVEC3F
        output y 0 0 0 # SFVEC3F
        output z 0 0 0 # SFVEC3F
    }
}

```

As the user drags the cursor over the ClickSensors' child geometry, the point on that geometry which lies directly beneath the cursor is determined. When isOver and isActive are TRUE, three output values are updated with information about the point on the object beneath the cursor. The output values are updated with information about the point on the surface of the object normal to the surface. The output values are updated with information about the point on the surface of the object normal to the surface. The ClickSensors' child geometry is mapped onto the surface, to support the 3D equivalent of an image map.

If the user presses the button associated with the pointing device while the cursor is located over its children, the ClickSensor will grab all further motion events from the pointing device until the button is released. The output press contains the last time the button was pressed over the ClickSensors. If the user presses the button associated with the pointing device while the cursor is located over its children, while the output click contains the last time the button was released after a grab by this ClickSensor. While the output click contains the last time the button was released after a grab by this ClickSensor, the output press remains grabbed, the output isActive remains TRUE. Motion of the pointing device while it has been grabbed by a ClickSensor is referred to as a "drag".

Issue: this is for locate-highlighting: Is that too much to ask, or should we assume fast picking from multiple locations?

The ClickSensor is triggered when the viewer points and clicks at its child geometry. The output enter contains the last time the cursor passed over any of the shape nodes contained in the ClickSensor, while the output exit contains the last time the cursor passed off of any of its contained shape nodes. While the cursor remains over any of the ClickSensor's children, the output isOver remains TRUE.

A PointingDeviceSensor is a node which tracks the pointing device input to its child geometry. A PointingDeviceSensor can be made inactive at any time by setting its enabled input to FALSE. An inactive sensor is reacted by scattering its enabled input to TRUE. There are two types of PointingDeviceSensors: ClickSensors, while the output remains over any of the ClickSensor's children, the output passes over any of the ClickSensor's children, the output isOver remains TRUE. The ClickSensor is triggered when the viewer points and clicks at its child geometry. The output enter contains the last time the cursor passed over any of the shape nodes contained in the ClickSensor, while the output exit contains the last time the cursor passed off of any of its contained shape nodes. While the cursor remains over any of the ClickSensor's children, the output isOver remains TRUE.

Pointing Device Sensors

```

FILE FORMAT/DEFAULTS
    PointProximitySensor {
        input enabled TRUE # SBOOL
        output center 0 # SFTIME
        output exit 0 # SFTIME
        output enter 0 # SFTIME
        output isIn 0 # SBOOL
        radius 0 # SFLOAT
        size 0 0 0 # SFVEC3F
        center 0 0 0 # SFVEC3F
        output isIn 0 0 0 # SFVEC3F
        output scroll 0 0 0 # SFVEC3F
        output touch 0 0 0 # SFVEC3F
        output volume 0 0 0 # SFVEC3F
        output x 0 0 0 # SFVEC3F
        output y 0 0 0 # SFVEC3F
        output z 0 0 0 # SFVEC3F
    }
}

```

The PointProximitySensor reports when the camera enters and leaves the sphere defined by the fields center and radius.

PointProximitySensor

```

FILE FORMAT/DEFAULTS
    VolumeProximitySensor {
        input enabled TRUE # SBOOL
        output center 0 # SFTIME
        output exit 0 # SFTIME
        output enter 0 # SFTIME
        output isIn 0 # SBOOL
        radius 0 # SFLOAT
        size 0 0 0 # SFVEC3F
        center 0 0 0 # SFVEC3F
        output isIn 0 0 0 # SFVEC3F
        output scroll 0 0 0 # SFVEC3F
        output touch 0 0 0 # SFVEC3F
        output volume 0 0 0 # SFVEC3F
        output x 0 0 0 # SFVEC3F
        output y 0 0 0 # SFVEC3F
        output z 0 0 0 # SFVEC3F
    }
}

```

A volumeproximity sensor uses volumes in the scene with have an active volume to use when that the world was entered, and can be used to start up animations or behaviors as soon as a world is loaded.

```

FILE FORMAT/DEFAULTS
discSensor {
    input enabled TRUE # SFB001
    output enter 0 # SFTIME
    output exit 0 # SFTIME
    output isOver FALSE # SFB001
    output press 0 # SFTIME
    output release 0 # SFTIME
    output rotate 0 # SFTIME
    output scroll 0 # SFTIME
}

```

The DiscSensor maps dragging motion into a rotation around the z axis of its local space. The rotation is as if you were scratching on a record turntable. The output rotation reflects the mapped drag position in local space.

DiscSensor

The output trackPoint always reflects the uncamped drag position in the x-y plane. The minPosition and maxPosition may be set to clamp the output translation to a range of values as measured from the origin of the x-y plane. If the x or y component of minPosition is less than or equal to the corresponding component of maxPosition, the output translation is not clamped in that dimension. The output trackPoint always reflects the uncamped drag position in the x-y plane.

```

FILE FORMAT/DEFAULTS
planeSensor {
    input enabled TRUE # SFB001
    output enter 0 # SFTIME
    output exit 0 # SFTIME
    output isOver FALSE # SFB001
    output press 0 # SFTIME
    output release 0 # SFTIME
    output scroll 0 # SFTIME
    output translate 0 0 0 # SFTIME
    output trackPoint 0 0 0 # SFTIME
    output translation 0 0 0 # SFTIME
    output trackPosition 0 0 0 # SFTIME
    output translate 0 0 0 # SFTIME
    output trackPosition 0 0 0 # SFTIME
    output trackTranslation 0 0 0 # SFTIME
}

```

The PlaneSensor maps dragging motion into a translation in two dimensions, in the x-y plane of its local space. The output translation reflects the mapped drag position in local space.

PlaneSensor

The output trackPoint always reflects the uncamped drag position along the x axis. Translation is not clamped. The output trackPoint always reflects the uncamped drag position along the x axis. If minPosition is less than or equal to maxPosition, the output translation is not clamped in the x axis. minPosition and maxPosition may be set to clamp the output translation to a range of values as measured from the origin of the x axis. If minPosition is less than or equal to maxPosition, the output translation is not clamped in the x axis. If minPosition is less than or equal to maxPosition, the output translation is not clamped in the x axis.

```

FILE FORMAT/DEFAULTS
lineSensor {
    input enabled TRUE # SFB001
    output enter 0 # SFTIME
    output exit 0 # SFTIME
    output isOver FALSE # SFB001
    output press 0 # SFTIME
    output release 0 # SFTIME
    output scroll 0 # SFTIME
    output translate 0 0 0 # SFTIME
    output trackPoint 0 0 0 # SFTIME
    output translation 0 0 0 # SFTIME
    output trackPosition 0 0 0 # SFTIME
    output trackTranslation 0 0 0 # SFTIME
}

```

The LineSensor maps dragging motion into a translation in one dimension, along the x axis of its local space. The output translation reflects the mapped drag position in local space.

LineSensor

DragScanners extend the ClickScanners' interaction. The outputs hitPoint, hitNormal, and hitTexture are active and implemented identically. The outputs hitPoint, hitNormal, and hitTexture are only updated upon the initial click down on the DragScanners' child geometry. There are five types of DragScanners: LineSensor and PlaneSensor support translation-oriented interfaces, and DiscSensor, ClickSensor and SphereSensor establish rotation-oriented interfaces.

A DragSensor tracks dragging over its child geometry just like the ClickSensor; however, raw screen position is mapped onto idealized 3D geometry. When the pointing device is pressed and dragged over the node's child geometry, the cursor's levers. Dragsensors track dragging in manner suitable for continuous controllers such as sliders, knobs, and buttons. When the pointing device is pressed and dragged over the node's child geometry, the cursors' raw screen position is mapped onto idealized 3D geometry.

The TimeSensor is a node which tracks the process of a desired time sequence. At the first stimulation tick until its input startTime is set to the beginning of a desired time sequence. A TimeSensor remains inactive

TimeSensor

Upon the initial click down on the DragSensors, child geometry, the hTipPoint determines the radius of the sphere used to map cursor input while dragging. The output trackPoint always reflects the uncaptured drag position on the face of this sphere, or in the plane perpendicular to the view vector if the cursor moves off of this sphere. The output always TRUE at the initial click down, and is set to FALSE if the cursor is dragged off of the sphere.

The free rotation of the SphereSensor is always uncamped.

```
FILE FORMAT/DEFUALTS
SphereSensor {
    Input enabled TRUE # SPP001
    Output center 0 # SPP111
    Output exit 0 # SPP111
    Output press 0 # SPP001
    Output isDowner FALSE # SPP001
    Output release 0 # SPP111
    Output rotate 0 0 1 0 # SPP001
    Output trackPoint 0 0 0 # SPP001
    Output onSphere 0 0 # SPP001
    Output hitTexture 0 0 # SPP001
    Output hitNormal 0 0 # SPP001
    Output isCaptured FALSE # SPP001
    Output isFree 0 # SPP001
    Output isSphere 0 0 # SPP001
    Output isCylinder 0 0 # SPP001
    Output isPlane 0 # SPP001
    Output isSphere 0 # SPP001
    Output isCylinder 0 # SPP001
    Output isPlane 0 # SPP001
}
```

The SphereSensor maps dragging motion into a free rotation about its center. The feel of the rotation is as if you were rolling a ball. The output rotation reflects the mapped drag position in local space.

SphereSensor

Upon the initial click down on the DragSensors, child geometry, the hTipPoint determines the radius of the cylinder used to map cursor input while dragging. The output trackPoint always reflects the uncaptured drag position on the face of this cylinder, or in the plane perpendicular to the view vector if the cursor moves off of this cylinder. The output always TRUE at the initial click down, and is set to FALSE if the cursor is dragged off of the cylinder.

The cylinder may be set to clamp the output rotation to a range of values as measured in radians about the y axis. If minAngle is less than or equal to maxAngle, the output rotation is not clamped.

```
FILE FORMAT/DEFUALTS
CylinderSensor {
    Input enabled TRUE # SPP001
    Output center 0 # SPP111
    Output exit 0 # SPP111
    Output press 0 # SPP001
    Output isDowner FALSE # SPP001
    Output release 0 # SPP111
    Output rotate 0 0 1 0 # SPP001
    Output trackPoint 0 0 0 # SPP001
    Output onCylinder 0 0 # SPP001
    Output hitTexture 0 0 # SPP001
    Output hitNormal 0 0 # SPP001
    Output isCaptured FALSE # SPP001
    Output isFree 0 # SPP001
    Output isSphere 0 # SPP001
    Output isCylinder 0 # SPP001
    Output isPlane 0 # SPP001
}
```

The CylinderSensor maps dragging motion into a rotation around the y axis of its local space. The feel of the rotation is as if you were turning a pin. The output rotation reflects the mapped drag position in local space.

CylinderSensor

The minAngle and maxAngle may be set to clamp the output rotation to a range of values as measured in radians about the z axis. If minAngle is less than or equal to maxAngle, the output rotation is not clamped. The output trackPoint always reflects the uncamped drag position in the x-y plane.

```
Output rotate 0 0 1 0 # SPP001
Output trackPoint 0 0 0 # SPP001
Output onSphere 0 0 # SPP001
Output hitTexture 0 0 # SPP001
Output hitNormal 0 0 # SPP001
Output isCaptured FALSE # SPP001
Output isFree 0 # SPP001
Output isSphere 0 # SPP001
Output isCylinder 0 # SPP001
Output isPlane 0 # SPP001
}
```

The description of outputting MF values belongs in the general interpolator section above, or maybe we should split up the interpolators into single-valued and multi-valued sections.

For example, if the first keyframe will be colors 0,1,2, the second colors 3,4,5, etc. The values are colors, the keys field must be an integer multiple of the number of keyframe times in the keys field; that integer multiple defines how many colors will be produced in the output. This node interpolates among a set of MFCOLOR values, to produce an MFCOLOR as output. The number of colors in the values field must be an integer multiple of the number of keyframe times in the keys field.

ColorInterpolator

```
FILE FORMAT/DEFAULTS
  interpolator {
    keys [ ] # MFCOLOR
    values [ ] # MFCOLOR
    input alpha 0 # SFRLOAT
    output outValue 0 # SFRLOAT
  }
```

This node interpolates among a set of SFRFLOAT values. The values field must contain exactly as many numbers as there are keyframe times in the keys field, or an error will be generated and results will be undefined.

FloatInterpolator

```
FILE FORMAT/DEFAULTS
  interpolator {
    keys [ ] # MFRFORMAT
    values [ ] # MFRFORMAT
    input alpha 0 # SFRFLOAT
    output outValue 1.000 # SFRFLOAT
  }
```

This node interpolates among a set of SFRFORMAT values. The values field must contain exactly as many rotations as there are keyframe times in the keys field, or an error will be generated and results will be undefined.

RotationInterpolator

Issue: do we want to do automatic SF/MF type conversion? Or define different interpolators?
We believe that keyframed animation will be common enough to justify the inclusion of these classes. We might choose to implement these as pre-defined prototypes of appropriate logically defined LOGIC nodes. Interpolators are nodes that are useful for doing keyframed animation. Given a sufficiently powerful scripting language, all of these interpolators could be implemented using LOGIC nodes (browsers) as built-in types.

Interpolators : animation

If cycleCount is <=0, the Timesensor will continue to tick continuously, without a cycle interval, in this case, cyclicalMethod and cycleMethod are ignored, and the output alpha remains at 0. This use of the Timesensor should be used with caution, since it incurrs continuous overhead on the simulation. In this case, cyclicalMethod outputs will resume to the current Time - pauseTime, outputs will resume hold on a stopwatch. It starts to the current Time - pauseTime, where they were interrupted, allowing the Timesensor to be used like the pause on a video deck.

The output pauseTime may be set to interrupt the progress of the Timesensor; at the first simulation tick whose real time >= pauseTime, outputs time and alpha resume on the original schedule. Following it to be used like the reset to startTime, outputs will resume to the current Time - pauseTime. If pauseTime is set to 0, the Timesensor will continue to tick continuously, without a cycle interval, in which they were interrupted, allowing the Timesensor to be used like the pause on a video deck.

```
FILE FORMAT/DEFAULTS
  timesensor {
    startTime 0 # SFTIME
    endTime 0 # SFTIME
    input pauseTime 0 # SFTIME
    input cycleInterval 0 # SFTIME
    output cycleCount 1 # SFTIME
    output cycleMethod FORWARD # SFTIME FORWARD | BACK | SWING
    output endTime 0 # SFTIME
    output alpha 0 # SFRFLOAT
  }
```

or it may be set to SWING, causing alpha to alternate 0 to 1, 1 to 0, on each successive interval. The input cycleCount may be set to repeat a fixed number of intervals, each of duration may be set to FORWARD, causing alpha between 0 and 1 over each interval. The input cycleMethod may be set to FORWARD, causing alpha to the same value between 0 and 1 over each interval, BACK which is equal to FORWARD.

where each unit is a single frame, the input cycleInterval may be set to drive continuous simulation or simulated behaviors. The input cycleCount to the input of other nodes to drive continuous simulation or simulated behaviors. The input cycleCount may be set to repeat a fixed number of intervals, each of duration may be set to FORWARD.

```
diffuseColor = LogIC { fields[script/inputs/etc...].color }
```

in the connection to diffuseColor.
this example 2D bytes smaller by not bothering to DEF/USE the Logic node, but instead just putting it
Also, the Logic node is defined to be a child of the ClickSensor for readability; we could have made

```
)  
color2 0 1 0  
color1 1 0 0  
isBeingClicked = USE CLICKSENSOR.lactive  
input SPCol1 isBeingClicked, output SPCol2 color1  
else if (isBeingClicked) color = color1; else color = color2;  
DEFLogicLogic ( inputs SPCol2 color2, inputs SPCol1 color1 )  
elements [ DEFLogicLogic ]
```

The Logic node was defined as the first child of the ClickSensor. The colors it produces are
hard-wired in the script, but could also have been provided as inputs:

```
)  
cube ()  
diffuseColor = USE LOGIC.color  
material {  
    isBeingClicked = USE CLICKSENSOR.lactive;  
    color = Color(0,1,0); // Green  
} else {  
    color = Color(1,0,0); // Red  
}  
DEFLogicLogic ( inputs SPCol1 isBeingClicked, output SPCol2 color )  
elements [ DEFLogicLogic ]
```

A cube that changes color when picked

All of these examples are "pseudo-VRML" to save space (coordinates are not specified, etc). They
also use a C++-like language for all Logic nodes.

Examples

Need to add way to interpolate value of a Switch.

```
)  
PILL FORMAT/DEFAULTS  
    VectorInterpolator {  
        keys 0 # MFFloat  
        values 0 # MFVec3f  
        output SFR10at  
        input alpfa 0 # MFVec3f  
        output outValue 0 # SFVec3f  
    }  
    output outValue 0 # SFVec3f
```

This node interpolates among a set of SFVec3f values, suitable for transforming normal vectors.

VectorInterpolator

```
)  
PILL FORMAT/DEFAULTS  
    PointInterpolator {  
        keys 0 # MFFloat  
        values 0 # MFVec3f  
        output SFR10at  
        input alpfa 0 # MFVec3f  
        output outValue 0 # SFVec3f  
    }  
    output outValue 0 # SFVec3f
```

This node interpolates among a set of SFVec3f values, doing linear interpolation. This would be
appropriate for interpolating a translation or a transformation.

PointInterpolator

```
)  
PILL FORMAT/DEFAULTS  
    PlaneInterpolator {  
        keys 0 # MFFloat  
        values 0 # MFColor  
        output SFR10at  
        input alpfa 0 # MFColor  
        output outValue 0 # MFColor  
    }  
    output outValue 0 # MFColor
```

```
PROTO HSVMaterial
    [ Input ClickCollector ambientColor is CONVERT_AMBIENT,
      Input ClickCollector diffuseColor is DIFFUSE ]
```

This prototype demonstrates why considering properties first-class objects is powerful. An HSV equivalent of the Material node is defined:

HSV color space Material node

This is overly simple; a real implementation would also have a Switch node that changes the button's appearance depending on its state. A good implementation would also use the isOver output of the ClickSensor to locate/highlight the button.

```
    ... Later, to use a toggle buttons:
        )
```

```
        ) ... Geometry of button is here ...
        )
```

```
        secret "if (toggle) state = !state; output = state"
    toggle = USE CLICKSENSOR.state
```

```
    toggle | input SPSOutput toggle, SPSOutput state, SPSOutput output ]
```

```
    DEFLOGIC logic ( DEFLOGICENSOR.clickSensor )
        output ison is LOGIC.outport ]
```

```
PROTO ToggleButton [ Eitled intialState is LOGIC.state,
    output ison is LOGIC.outport ]
```

This is a very simple implementation of a button that changes its output when clicked:

A simple toggle-button prototype

A Logic node could also be inserted between the Transform and the TimeSensor (for example) start arbitrary scheduling can be done using Logic nodes with multiple inputs/outputs.

```
seconds, animation #2 starts when animation #2 ends (19 seconds from now), etc-
```

```
seconds, respectively, animation #2 and #3 start 14 seconds from now and last 5 and 7
```

```
away and lasts 14 seconds, animation #1 starts when animation #1 starts right
```

```
pressed or a puzzle has been solved, or to start a series of animations (e.g., animation #1 starts right
```

```
the animation 3 seconds after the cube is picked, or to only start an animation If a button has been
```

```
pressed for 10 seconds, animation #1 starts when animation #1 starts right
```

```
and rotate the objects, we could just add a set of keyframes to the Transform's rotation field;
```

when the user clicks and releases the mouse (or other pointing device) over the cube. To both animate

this example just starts a 10-second keyframe animation that changes the translation of some objects

```
    ) ... objects to be animated ...
```

```
    ) .outValue
    ) alpfa = USE TIMESENSOR.alpfa
```

```
    ) values [ ... ] keys [ ... ] .outValue
```

```
    ) translation = PointInterpolator ( ... ) .outValue
```

```
    ) transform (
```

```
    ) .startTime = USE START_RELATIVETIME
```

```
    ) interval 10.0
    ) released from the cube:
    # animate for 10 seconds starting whenever the button is
```

```
# for readability, I put the timesensor here:
    DEFTIMESENSOR timesensor (
```

```
    ) separator {
```

```
    # And this is the object that will move:
    ) cube ( )
```

```
    ) DEFSSTART_CLICKSENSOR (
```

```
    # This is the cube that will start the animation:
```

```
    ) separator {
```

Start a keyframe animation whenever a cube is picked

So, in our example, the "...AvatarGeometry..." would actually simply be:

The script would ask the server (given in serviceName) for the geometry for some object (whose name is given in objectName). If this object might change over time, the script will start an asynchronous process that changes the geometry whenever it gets a change message from the server.

*...geometric
script ...
objectName ... serviceName ...
fields [setting objectName, setting serviceName,
nodeReference = DEF L LogIC
field setting objectName IS L.objectName,
PROTO FromServer [field setting objectName IS L.objectName,
nodeReference [field setting objectName IS L.objectName,*

The second part of the problem is displaying the avatars. To make this easier to read, let's define a "FromServer" prototype that lets us refer to a node that we get from the server.

So, first, the task of reporting "my" position to the server is accomplished by using a VolumeProximitySensor that is big enough to surround the entire world, with my current position inside that world-sized region being reported to the server by using a Logic node. The Logic script just has an ImpulseChange method that reads the current position and sends it to the server (it's all assuming that your language supports that kind of communication). Details are left as an exercise for the hacker.

*...AvatarGeometry
script ...
input currentPosition = USE positionReporter.currentPosition
logic [
DEF positionReporter VolumeProximitySensor (
separater (... geometry of the world ...)
#VRML V1.0 utf8*

There are two main parts to this example: Reporting "our" position to the Avatar server, and getting Avatars geometry and the positions of all Avatars from the server. The overall structure of the world looks like:

Talking to a multi-user Avatar motion server

Useful properties such as a camera defined by yaw/pitch/roll values could also be implemented this way.

Note that attaching a ColorInterpolator (which interpolates in RGB space) to an HSVMaterial gives color interpolation in HSV space (the "RGB" values that are being interpolated will be translated into HSV values as they're interpolated by the HSVMaterial's logic).

*ambientColor = USE CONVERT, diffuseColor
diffuseColor = USE CONVERT, diffuseColor
emissiveColor = USE CONVERT, emissiveColor
specularColor = USE CONVERT, specularColor
done.
set correspondence output
convert to HSV
for all values:
for each do:
read ambient/diffuse/specular/emissive:
script, a inputchanged method does:
etc ...
ambientIn .2 .2 .2 # Define default value
etc ...
ambientColor = DEF CONVERT logIC {
DEF N Material {
input MECOLOR ambientIn;
output MECOLOR ambientOut;
ambientColor = DEF CONVERT logIC {
implementation: Material hooked up to logic:
input HPRotate transparency IS M.transparency;
input MCOLOR emissiveColor IS CONVERT, emissiveIn;
input MCOLOR specularColor IS CONVERT, specularIn;*

When an Avatar moves, the server might tell the script to just modify some transformation in the AllAvatars scene graph. That will work well for browsers that don't optimize the scene and that use most efficiently for those other browsers, the AllAvatars scene graph could itself contain FromServer the scene graph as their internal representation; it won't work very well for other browsers. To be NodeReference objects, like this:

```
serverName "...some protocol/address..."  
# ALLAvatars scene created by script (from server info):  
Separator {  
    # Avatar 1  
    # Avatar 2  
    ... geometry for avatar 1  
}  
fromServer {  
    objectName "Avatar2Transform"  
    serverName "...http://www.sgi.com/Q_14_12_server."  
}  
fromServer {  
    separator {  
        # Avatar 1  
        # Avatar 2  
        ... etc, for each avatar in this world  
    }  
}
```