

Australian Standard™

**Information technology—Coding of
audio-visual objects**

**Part 16: Animation Framework
eXtension (AFX)**



This Australian Standard was prepared by Committee IT-029, Coded Representation of Picture, Audio and Multimedia/Hypermedia Information. It was approved on behalf of the Council of Standards Australia on 27 October 2004.
This Standard was published on 3 December 2004.

The following are represented on Committee IT-029:

Australian Broadcasting Corporation
Australian Consumer Association
Australian Subscription Television
CSIRO Information and Communication Technology Centre
DSTC
Department of Defence (Australia)
Free TV Australia
Special Broadcasting Service
The University of New South Wales
University of Sydney Australian Broadcasting Authority
University of Wollongong
Victoria University of Technology

Keeping Standards up-to-date

Standards are living documents which reflect progress in science, technology and systems. To maintain their currency, all Standards are periodically reviewed, and new editions are published. Between editions, amendments may be issued. Standards may also be withdrawn. It is important that readers assure themselves they are using a current Standard, which should include any amendments which may have been published since the Standard was purchased.

Detailed information about Standards can be found by visiting the Standards Web Shop at www.standards.com.au and looking up the relevant Standard in the on-line catalogue.

Alternatively, the printed Catalogue provides information current at 1 January each year, and the monthly magazine, *The Global Standard*, has a full listing of revisions and amendments published each month.

Australian Standards™ and other products and services developed by Standards Australia are published and distributed under contract by SAI Global, which operates the Standards Web Shop.

We also welcome suggestions for improvement in our Standards, and especially encourage readers to notify us immediately of any apparent inaccuracies or ambiguities. Contact us via email at mail@standards.org.au, or write to the Chief Executive, Standards Australia International Ltd, GPO Box 5420, Sydney, NSW 2001.

This Standard was issued in draft form for comment as DR 04417.

Australian Standard™

Information technology—Coding of audio-visual objects

Part 16: Animation Framework eXtension (AFX)

First published as AS ISO/IEC 14496.16—2004.

COPYRIGHT

© Standards Australia International

All rights are reserved. No part of this work may be reproduced or copied in any form or by any means, electronic or mechanical, including photocopying, without the written permission of the publisher.

Published by Standards Australia International Ltd GPO Box 5420, Sydney, NSW 2001, Australia

ISBN 0 7337 6370 7

PREFACE

This Standard was prepared by the Standards Australia Committee IT-029, Coded Representation of Picture, Audio and Multimedia/Hypermedia Information.

This Standard is identical with, and has been reproduced from, ISO/IEC 14496-16:2004, *Information technology—Coding of audio-visual objects, Part 16: Animation Framework eXtension (AFX)*.

The objective of this Standard is to provide the multimedia designer with an understanding of models and their relationships for interactive multi-media contents including (but not limited to) computer games and animation. This part also provides the designer with information on progressive streaming of view-dependent information.

The terms ‘normative’ and ‘informative’ are used to define the application of the annex to which they apply. A normative annex is an integral part of a standard, whereas an informative annex is only for information and guidance.

As this Standard is reproduced from an international standard, the following applies:

- (a) Its number appears on the cover and title page while the international standard number appears only on the cover.
- (b) In the source text ‘this International Standard’ should read ‘this Australian Standard’.
- (c) A full point substitutes for a comma when referring to a decimal marker.

References to International Standards should be replaced by references to Australian or Australian/New Zealand Standards, as follows:

<i>Reference to International Standard</i>		<i>Australian Standard/New Zealand</i>	
ISO/IEC		AS/NZS	
14496	Information technology—Coding of audio-visual objects	14496	Information technology—Coding of audio-visual objects
14496-1	Part 1: Systems	14496.1	Part 1: Systems
14496-2	Part 2: Visual	14496.2	Part 2: Visual

Only referenced documents that have been adopted as Australian or Australian/New Zealand Standards have been listed.

CONTENTS

Page

1	Scope	1
2	Normative references	1
3	Symbols and abbreviated terms	1
4	Animation Framework eXtension (AFX)	2
4.1	Introduction	2
4.2	The AFX model	3
4.3	Geometry tools	5
4.3.1	Introduction	5
4.3.2	Non-Uniform Rational B-Spline (NURBS)	6
4.3.3	Subdivision surfaces	9
4.3.4	MeshGrid representation	27
4.3.5	Particle systems	33
4.4	Solid representation	37
4.4.1	Solid modeling nodes	37
4.4.2	Scripting extensions for solid modeling operations	41
4.4.3	Examples	45
4.5	Texture tools	47
4.5.1	Introduction	47
4.5.2	Depth Image-Based Representation	47
4.5.3	Light-field mapping	52
4.6	Modeling tools	57
4.6.1	NonLinearDeformer	57
4.6.2	Free-form deformations	58
4.7	Generic skeleton, muscle and skin-based model definition and animation	59
4.7.1	Generic skeleton, muscle and skin-based model definition	59
4.7.2	Generic skeleton, muscle and skin-based model animation	65
5	Bitstream specification	67
5.1	Wavelet Subdivision Surfaces	67
5.1.1	Downstream syntax	67
5.1.2	Upstream syntax (for backchannel)	69
5.2	MeshGrid stream	70
5.2.1	Overview	70
5.2.2	The Bitstream structure	70
5.2.3	Decoder	86
5.3	Depth Image-Based Representation	97
5.3.1	Octree Compression	97
5.4	Compressed Bone-based animation	102
5.4.1	Bone-Based Animation stream encapsulation within BIFS-Anim	102
5.4.2	Bone-Based Animation stream definition	103
5.5	AFX Generic Backchannel	115
5.5.1	Usage of the backchannel	115
5.5.2	BackChannel Access Unit	115
5.5.3	BackChannelCommand	115
Annex A	(normative) Wavelet Mesh Decoding Process	127
A.1	Overview	127
A.2	Base mesh	127
A.3	Definitions and notations	127
A.4	Bitplanes extraction	128
A.5	Zero-tree decoder	128
A.6	Synthesis filters and mesh reconstruction	128

A.7	Basis change.....	130
Annex B	(normative) MeshGrid Representation.....	131
B.1	The hierarchical multi-resolution MeshGrid	131
B.1.1	Relation between the resolution level of the mesh and the number of reference-surfaces.....	132
B.1.2	Relation between ROIs and Tiles.....	133
B.1.3	Relation between the resolution level of the grid and the number of reference-surfaces.	133
B.2	Scalability Modes.....	134
B.2.1	The mesh resolution scalability.....	134
B.2.2	The scalability in shape precision	134
B.2.3	The vertex position scalability	135
B.3	Animation Possibilities.....	136
B.3.1	Vertex-based animation.....	136
B.3.2	Rippling effects, changing the offsets	137
B.3.3	Animation of the reference-grid points	137
Annex C	(informative) MeshGrid representation.....	139
C.1	Representing Quadrilateral meshes in the MeshGrid format	139
C.2	IndexedFaceSet models represented by the MeshGrid format.	140
Annex D	(informative) Solid representation	142
D.1	Overview	142
D.2	Solid Primitives	142
D.3	Arithmetics of Forms	143
D.3.1	General Arithmetic Operators on Densities	144
D.3.2	Ternary Logic : The Kleene Operators	146
D.3.3	Densities Filtering	148
Annex E	(normative) Light-field mapping framework	149
Annex F	(informative) Generic skeleton, muscle and skin-based model definition and animation: XMT compliant animation file format	154
F.1	"bone" animation mask and animation values	154
F.2	"muscle" animation mask and animation values	154
F.3	Frame mask and frame values representation	154
F.4	XSD for BBA file.....	155
Annex G	(informative) Face and Body animation: XMT compliant animation and encoding parameter file format.....	159
G.1	XSD for FAP file format	159
G.2	XSD for BAP file format	160
G.3	XSD for EPF.....	162
Annex H	(normative) Node coding tables	165
Annex I	(informative) Patent statements	183
Bibliography		184

Information technology — Coding of audio-visual objects —

Part 16:

Animation Framework eXtension (AFX)

1 Scope

This International Standard specifies MPEG-4 Animation Framework eXtension (AFX) model for creating interactive multimedia contents by composing natural and synthetic objects. Within this model, MPEG-4 is extended with higher-level synthetic objects for geometry, texture, and animation as well as dedicated compressed representations.

AFX also specifies a backchannel for progressive streaming of view-dependent information.

2 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 14496-1:2001, *Information technology — Coding of audio-visual objects — Part 1: Systems*

ISO/IEC 14496-2, *Information technology — Coding of audio-visual objects — Part 2: Visual*

ISO/IEC 14496-11:2003, *Information technology — Coding of audio-visual objects — Part 11: Scene description and application engine*

3 Symbols and abbreviated terms

List of symbols and abbreviated terms.

AFX	Animation Framework eXtension
BIFS	Binary Format for Scene
DIBR	Depth-Image Based Representation
ES	Elementary Stream
IBR	Image-Based Rendering
LFM	Light-Field Mapping
NDT	Node Data Type
OD	Object Descriptor
SDL	Syntactic Description Language
VRML	Virtual Reality Modelling Language

4 Animation Framework eXtension (AFX)

4.1 Introduction

"Most people think the word 'animation' means movement. But it doesn't. It comes from 'animus' which means 'life or to live'. Making it move is not animation, but just the mechanics of it."

Frank Thomas and Ollie Johnston
Disney Animation: the illusion of life, 1981

MPEG-4, ISO/IEC 14496, is a multimedia standard that enables composition of multiple audio-visual objects on a terminal. Audio and visual objects can come from *natural* sources (e.g. a microphone, a camera) or from *synthetic* ones (i.e. made by a computer); each source is called a *media* or a *stream*. On their terminals, users can display, play, and interact with MPEG-4 audio-visual contents, which can be downloaded previously or streamed from remote servers. Moreover, each object may be protected to ensure a user has the right credentials before downloading and displaying it.

Unlike natural audio and video objects, computer graphics objects are purely synthetic. Mixing computer graphics objects with traditional audio and video enables augmented reality applications, i.e. applications mixing natural and synthetic objects. Examples of such contents range from DVD menus, and TV's Electronic Programming Guides to medical and training applications, games, and so on.

Like other computer graphics specifications, MPEG-4 synthetic objects are organized in a *scene graph* based on VRML97, which is a direct acyclic tree where *nodes* represent objects and branches their properties, called *fields*. As each object can receive and emit events, two branches can be connected by the means of a *route*, which propagates events from one field of one node to another field of another node. As any other MPEG-4 media, scenes may receive updates from a server that modify the topology of the scene graph.

The Animation Framework eXtension (AFX) proposes a conceptual organization of synthetic models for computer animations as well as specific compression schemes; models defined in ISO/IEC 14496 Part 11 and in this part fit in organization (see Figure 2).




Figure 1 — Animation Framework and MPEG-4 Systems

Figure 1 shows the position of the Animation Framework within MPEG-4 Terminal architecture. It extends the existing BIFS tree with new tools and define the AFX streams that carry dedicated compressed object representations and a backchannel for view-dependent features.

4.2 The AFX model

To understand the AFX model, let's take an example. Suppose one wants to build an avatar. The avatar consists of geometry elements that describe its legs, arms, head and so on. Simple geometric elements can be used and deformed to produce more physically realistic geometry. Then, skin, hair, cloths are added. These may be physic-based models attached to the geometry. Whenever the geometry is deformed, these models deform and thanks to their physics, they may produce wrinkles. Biomechanical models are used for motion, collision response, and so on. Finally, our avatar may exhibit special behaviors when it encounters objects in its world. It might also learn from experiences: for example, if it touches a hot surface and gets hurt, next time, it will avoid touching it.

This hierarchy also works in a top to bottom manner: if it touches a hot surface, its behavior may be to retract its hand. Retracting its hand follows a biomechanical pattern. The speed of the movement is based on the physical property of its hand linked to the rest of its body, which in turn modify geometric properties that define the hand.

AFX defines 6 groups of components, following [39].




Figure 2 — Models in computer games and animation [39]

1. *Geometric component.* These models capture the form and appearance of an object. Many characters in animations and games can be quite efficiently controlled at this low-level. Due to the predictable nature of motion, building higher-level models for characters that are controlled at the geometric level is generally much simpler.
2. *Modeling component.* These models are an extension of geometric models and add linear and non-linear deformations to them. They capture transformation of the models without changing its original shape. Animations can be made on changing the deformation parameters independently of the geometric models.
3. *Physical component.* These models capture additional aspects of the world such as an object's mass inertia, and how it responds to forces such as gravity. The use of physical models allows many motions to be created automatically and with unparalleled realism.
4. *Biomechanical component.* Real animals have muscles that they use to exert forces and torques on their own bodies. If we already have built physical models of characters, they can use virtual muscles to move themselves around. These models have their roots in control theory.
5. *Behavioral component.* A character may expose a *reactive behavior* when its behavior is solely based on its perception of the current situation (i.e. no memory of previous situations). *Goal-directed behaviors* can be used to define a cognitive character's goals. They can also be used to model *flocking behaviors*.
6. *Cognitive component.* If the character is able to learn from stimuli from the world, it may be able to adapt its behavior. These models are related to artificial intelligence techniques.

AFX specification currently deals with the first four categories. Models of the last two categories are typically application-specific and often designed programmatically. Similarly, while the first four categories can be animated using existing tools such as interpolators, the last two categories have their own logic and cannot be animated the same way.

In each category, one can find many models for any applications: from simple models that require little processing power (low-level models) to more complex models (high-level models) that require more computations by the terminal. VRML [4] and BIFS [3] specifications provide low-level models that belong to the Geometry component with the exception of Face and Body Animation tools that belong to the Biomechanical component.

Higher-level components can be defined as providing a compact representation of functionality in a more abstract manner. Typically, this abstraction leads to mathematical models that need few parameters. These models cannot

be rendered directly by a graphic card: internally, they are converted to low-level primitives a graphic card can render. Besides a more compact representation, this abstraction often provides other functionalities such as but not limited to compact representation, view-dependent subdivision, automatic level-of-details, smoother graphical representation, scalability across terminals, and progressive local refinements.

For example, a subdivision surface can be subdivided based on the area viewed by the user. For animations, piecewise-linear interpolators require few computations but require lots of data in order to represent a curved path. Higher-level animation models represent animation using piecewise-spline interpolators with less values and provide more control over the animation path and timeline.

In the remaining of this document, this conceptual organization of tools is followed in the same spirit an author will create content: the geometry is first defined with or without solid modeling tools, then texture is added to it. Objects can be deformed and animated using modeling and animation tools. Finally, avatars need biomechanical tools. Behavioral and Cognitive models can be programmatically implemented using JavaScript or MPEG-J defined in ISO/IEC 14496-11 [3].

NOTE

Some generic tools developed originally within the Animation Framework eXtension have been relocated in ISO/IEC 14496-11/AMD1 [3] along with other generic tools. This includes:

- Spline-based generic animation tools, called Animator nodes;
- Optimized interpolator compression tools;
- BitWrapper node that enables compressed representation of existing nodes;
- Procedural textures based on fractal plasma.

4.3 Geometry tools

4.3.1 Introduction

Geometry tools consist of the following technologies:

- NURBS, which consists of the following nodes: **NurbsSurface**, **NurbsCurve**, and **NurbsCurve2D**;
- Subdivision surfaces consisting of the following nodes: **SubdivisionSurface**, **SubdivSurfaceSector** and **WaveletSubdivisionSurface** of which the latter enables to add wavelet-encoded details at different resolutions to subdivision surfaces;
- MeshGrid, which consists of the **MeshGrid** node;
- Particle systems, which consists of the following nodes: **Particles**, **ParticleInitBox**, **ParticleInitCone**, **ParticleObstacle**, and **PointAttractor**.

4.3.2 Non-Uniform Rational B-Spline (NURBS)

4.3.2.1 Introduction

A Non-Uniform Rational B-Spline (NURBS) curve of degree $p > 0$ (and hence of order $p+1$) and control points $\{\mathbf{P}_i\}$ ($0 \leq i \leq n-1$; $n \geq p+1$) has for Equation 1 [39], [66], [89]:

$$\mathbf{C}(u) = \frac{\sum_{i=0}^{n-1} R_{i,p}(u) \mathbf{P}_i}{\sum_{i=0}^{n-1} N_{i,p}(u) w_i} = \frac{\sum_{i=0}^{n-1} N_{i,p}(u) w_i \mathbf{P}_i}{\sum_{i=0}^{n-1} N_{i,p}(u) w_i}$$

Equation 1 — NURBS curve definition

The parameter $u \in [0, 1]$ allows to travel along the curve and $\{R_{i,p}\}$ are its rational basis functions. The latter can in turn be expressed in terms of some positive (and not all null) weights $\{w_i\}$, and $\{N_{i,p}\}$, the p^{th} -degree B-Spline basis functions defined on the possibly non-uniform, but always non-decreasing knot sequence/vector of length $m = n + p + 1$: $U = \{u_0, u_1, \dots, u_{m-1}\}$, where $0 \leq u_i \leq u_{i+1} \leq 1 \ \forall \ 0 \leq i \leq m-2$.

The B-Spline basis functions are defined recursively using the Cox de Boor formula:

$$N_{i,0}(u) = \begin{cases} 1 & \text{if } u_i \leq u < u_{i+1}; \\ 0 & \text{otherwise;} \end{cases}$$

$$N_{i,p}(u) = \frac{u - u_i}{u_{i+p} - u_i} N_{i,p-1}(u) + \frac{u_{i+p+1} - u}{u_{i+p+1} - u_{i+1}} N_{i+1,p-1}(u).$$

If $u_i = u_{i+1} = \dots = u_{i+k-1}$, it is said that u_i has multiplicity k . $\mathbf{C}(u)$ is infinitely differentiable inside knot spans, *i.e.*, for $u_i < u < u_{i+1}$, but only $p-k$ times differentiable at the parameter value corresponding to a knot of multiplicity k , so setting several consecutive knots to the same value u_i decreases the smoothness of the curve at u_i . In general, knots cannot have multiplicity greater than p , but the first and/or last knot of U can have multiplicity $p+1$, *i.e.*, $u_0 = \dots = u_p = 0$ and/or $u_{m-p-1} = \dots = u_{m-1} = 1$, which causes $\mathbf{C}$ to interpolate the corresponding endpoint(s) of the control polygon defined by $\{\mathbf{P}_i\}$, *i.e.*, $\mathbf{C}(0) = \mathbf{P}_0$ and/or $\mathbf{C}(1) = \mathbf{P}_{n-1}$. Therefore, a knot vector of the kind

$U = \left\{ \underbrace{u_0 = 0, \dots, 0}_{p+1}, u_{p+1}, \dots, u_{m-p-2}, \underbrace{1, \dots, 1}_{p+1} = u_{m-1} \right\}$ causes the curve to be endpoint interpolating, *i.e.*, to interpolate both

endpoints of its control polygon. Extreme knots, multiple or not, may enclose any non-decreasing subsequence of interior knots: $0 < u_i \leq u_{i+1} < 1$. An endpoint interpolating curve with no interior knots, *i.e.*, one with U consisting of $p+1$ zeroes followed by $p+1$ ones, with no other values in between, is a p^{th} -degree Bézier curve: *e.g.*, a cubic Bézier curve can be described with four control points (of which the first and last will lie on the curve) and a knot vector $U = \{0, 0, 0, 0, 1, 1, 1, 1\}$.

It is possible to represent all types of curves with NURBS and, in particular, all conic curves (including parabolas, hyperbolas, ellipses, etc.) can be represented using rational functions, unlike when using merely polynomial functions.

Other interesting properties of NURBS curves are the following:

- Affine invariance: rotations, translations, scalings, and shears can be applied to the curve by applying them to $\{\mathbf{P}_i\}$.
- Convex hull property: the curve lies within the convex hull defined by $\{\mathbf{P}_i\}$. The control polygon defined by $\{\mathbf{P}_i\}$ represents a piecewise approximation to the curve. As a general rule, the lower the degree, the closer the NURBS curve follows its control polygon.

- Local control: if the control point $\mathbf{P}_i$ is moved or the weight w_i is changed, only the portion of the curve swept when $u_i < u < u_{i+p+1}$ is affected by the change.

NURBS surfaces are defined as tensor products of two NURBS curves, of possibly different degrees and/or numbers of control points:

$$C(u, v) = \frac{\sum_{i=0}^{n-1} \sum_{j=0}^{l-1} N_{i,p}(u) N_{j,q}(v) w_{i,j} \mathbf{P}_{i,j}}{\sum_{i=0}^{n-1} \sum_{j=0}^{l-1} N_{i,p}(u) N_{j,q}(v) w_{i,j}}$$

Equation 2 — NURBS surface definition

The two independent parameters $u, v \in [0, 1]$ allow to travel across the surface. The B-Spline basis functions are defined as previously, and the resulting surface has the same interesting properties that NURBS curves have. Multiplicity of knots may be used to introduce sharp features (corners, creases, etc.) in an otherwise smooth surface, or to have it interpolate the perimeter of its control polyhedron.

4.3.2.2 NurbsCurve

4.3.2.2.1 Node interface

NurbsCurve { #%NDT=SFGeometryNode

eventIn	MFInt32	set_colorIndex	
exposedField	SFColorNode	color	NULL
exposedField	MFVec4f	controlPoint	[]
exposedField	SFInt32	tessellation	0 # [0, ∞)
field	MFInt32	colorIndex	[] # [-1, ∞)
field	SFBool	colorPerVertex	TRUE
field	MFFloat	knot	[] # [0, 1]
field	SFInt32	order	4 # [3, 34]

}

4.3.2.2.2 Functionality and semantics

The **NurbsCurve** node describes a 3D NURBS curve, which is displayed as a curved line, similarly to what is done with the **IndexedLineSet** primitive.

The **order** field defines the order of the NURBS curve, which is its degree plus one.

The **controlPoint** field defines a set of control points in a coordinate system where the weight is the last component. The number of control points must be greater than or equal to the order of the curve. All weight values must be greater than or equal to 0, and at least one weight must be strictly greater than 0. If the weight of a control point is increased above 1, that point is more closely approximated by the surface. However the surface is not changed if all weights are multiplied by a common factor.

The **knot** field defines the knot vector. The number of knots must be equal to the number of control points plus the order of the curve, and they must be ordered non-decreasingly. By setting consecutive knots to the same value, the degree of continuity of the curve at that parameter value is decreased. If o is the value of the field **order**, o consecutive knots with the same value at the beginning (resp. end) of the knot vector cause the curve to interpolate the first (resp. last) control point. Other than at its extremes, there may not be more than $o-1$ consecutive knots of equal value within the knot vector. If the length of the knot vector is 0, a default knot vector consisting of o 0's followed by o 1's, with no other values in between, will be used, and a Bézier curve of degree $o-1$ will be obtained. A closed curve may be specified by repeating the starting control point at the end and specifying a periodic knot vector.

The **tessellation** field gives hints to the curve tessellator as to the number of subdivision steps that must be used to approximate the curve with linear segments: for instance, if the value t of this field is greater than or equal to that of the **order** field, t can be interpreted as the absolute number of tessellation steps, whereas $t = 0$ lets the browser choose a suitable tessellation.

Fields **color**, **colorIndex**, **colorPerVertex**, and **set_colorIndex** have the same semantic as for **IndexedLineSet** applied to the control points.

4.3.2.3 NurbsCurve2D

4.3.2.3.1 Node interface

```
NurbsCurve2D { #NDR=SFGeometryNode
  eventIn      MFInt32      set_colorIndex
  exposedField SFCOLORNode  color      NULL
  exposedField MFVec3f      controlPoint []
  exposedField SFInt32      tessellation 0 # [0, ∞)
  field        MFInt32      colorIndex   [] # [-1, ∞)
  field        SFBBool      colorPerVertex TRUE
  field        MFFloat      knot         [] # [0, 1]
  field        SFInt32      order        4 # [3, 34]
}
```

4.3.2.3.2 Functionality and semantics

The **NurbsCurve2D** is the 2D version of **NurbsCurve**; it follows the same semantic as **NurbsCurve** with 2D control points.

4.3.2.4 NurbsSurface

4.3.2.4.1 Node interface

```
NurbsSurface { #NDR=SFGeometryNode
  eventIn      MFInt32      set_colorIndex
  eventIn      MFInt32      set_texCoordIndex
  exposedField SFCOLORNode  color      NULL
  exposedField MFVec4f      controlPoint []
  exposedField SFTextureCoordinateNode texCoord NULL
  exposedField SFInt32      uTessellation 0 # [0, ∞)
  exposedField SFInt32      vTessellation 0 # [0, ∞)
  field        MFInt32      colorIndex   [] # [-1, ∞)
  field        SFBBool      colorPerVertex TRUE
  field        SFBBool      solid        TRUE
  field        MFInt32      texCoordIndex [] # [-1, ∞)
  field        SFInt32      uDimension   4 # [3, 258]
  field        MFFloat      uKnot        [] # (-∞, ∞)
  field        SFInt32      uOrder       4 # [3, 34]
  field        SFInt32      vDimension   4 # [3, 258]
  field        MFFloat      vKnot        [] # (-∞, ∞)
  field        SFInt32      vOrder       4 # [3, 34]
}
```

4.3.2.4.2 Functionality and semantics

The **NurbsSurface** Node describes a 3D NURBS surface. Similar 3D surface nodes are **ElevationGrid** and **IndexedFaceSet**. In particular, **NurbsSurface** extends the definition of **IndexedFaceSet**. If an

implementation does not support **NurbsSurface**, it should still be able to display its control polygon as a set of (triangulated) quadrilaterals, which is the coarsest approximation of the NURBS surface.

The **uOrder** field defines the order of the surface in the *u* dimension, which is its degree in the *u* dimension plus one. Similarly, the **vOrder** field define the order of the surface in the *v* dimension.

The **uDimension** and **vDimension** fields define respectively the number of control points in the *u* and *v* dimensions, which must be greater than or equal to the respective orders of the curve.

The **controlPoint** field defines a set of control points in a coordinate system where the weight is the last component. These control points form an array of dimension **uDimension** x **vDimension**, in a similar way to the regular grid formed by the control points of an **ElevationGrid**, the difference being that, in the case of a **NurbsSurface**, the points need not be regularly spaced. The number of control points in each dimension must be greater than or equal to the corresponding order of the curve. Specifically, the three spatial coordinates (*x*, *y*, *z*) and weight (*w*) of control point P_{ij} (NB: $0 \leq i < \text{uDimension} \geq \text{uOrder}$ and $0 \leq j < \text{vDimension} \geq \text{vOrder}$) are obtained as follows:

```
P[i,j].x = controlPoint[i + j*uDimension].x
```

```
P[i,j].y = controlPoint[i + j*uDimension].y
```

```
P[i,j].z = controlPoint[i + j*uDimension].z
```

```
P[i,j].w = controlPoint[i + j*uDimension].w
```

All **weight** values must be greater than or equal to 0, and at least one weight must be strictly greater than 0. If the weight of a control point is increased above 1, that point is more closely approximated by the surface. However the surface is not changed if all weights are multiplied by a common factor.

The **uKnot** and **vKnot** fields define the knot vectors in the *u* and *v* dimensions respectively, and their semantics are analogous to that of the **knot** field of the **NurbsCurve** node.

The **uTessellation** and **vTessellation** fields give hints to the surface tessellator as to the number of subdivision steps that must be used to approximate the surface with (triangulated) quadrilaterals: for instance, if the value *t* of the **uTessellation** field is greater than or equal to that of the **uOrder** field, *t* can be interpreted as the absolute number of tessellation steps in the *u* dimension, whereas *t* = 0 lets the browser choose a suitable tessellation.

solid is a rendering hints that indicates if the surface is closed. If FALSE, two-side lighting is enabled. If TRUE, only one-side lighting is applied.

color, **colorIndex**, **colorPerVertex**, **set_colorIndex**, **texCoord**, **texCoordIndex**, and **set_textCoordIndex** have the same functionality and semantics as for **IndexedFaceSet** and they apply on the control points defined in **controlPoint**, just like these fields apply on **coord** points of an **IndexedFaceSet**.

4.3.3 Subdivision surfaces

4.3.3.1 Introduction

Subdivision is a recursive refinement process that splits the facets or vertices of a polygonal mesh (the initial "control hull") to yield a smooth limit surface. The refined mesh obtained after each subdivision step is used as the control hull for the next step, so all successive (and hierarchically nested) meshes can be regarded as control hulls. The refinement of a mesh is performed both on its topology, as the vertex connectivity is made richer and richer, and on its geometry, as the new vertices are positioned in such a way that the angles formed by the new facets are smaller than those formed by the old facets.




Figure 3 — Nested control meshes

Figure 3 shows an example of how subdivision would be applied to a triangular mesh. At each step, one triangle is shaded to highlight the correspondence between one level and the next: each triangle is broken down into four new triangles, whose vertex positions are perturbed to have a smoother and smoother surface.

4.3.3.1.1 Piecewise Smooth Surfaces

Subdivision schemes for piecewise smooth surfaces include common modeling primitives such as quadrilateral free-form patches with creases and corners.

A somewhat informal description of piecewise smooth surfaces is given here; for simplicity, only surfaces without self intersections are considered. For a closed C^1 -continuous surface in $\mathfrak{R}^3$, each point has a neighborhood that can be smoothly deformed into an open planar disk D . A surface with a *smooth boundary* can be described in a similar way, but neighborhoods of boundary points can be smoothly deformed into a half-disk H , with closed boundary (Figure 4). In order to allow *piecewise* smooth boundaries, two additional types of local charts are introduced: concave and convex corner charts, Q_3 and Q_1 .




Figure 4 — The charts for a surface with piecewise smooth boundary

Piecewise-smooth surfaces are constructed out of surfaces with piecewise smooth boundaries joined together. If two surface patches have a common boundary, but different normal directions along the boundary, the resulting surface has a sharp crease.

Two adjacent smooth segments of a boundary are allowed to be joined, producing a crease ending in a *dart* (cf. [44]). For dart vertices an additional chart Q_0 is required; the surface near a dart can be deformed into this chart smoothly everywhere except at an open edge starting at the center of the disk.

4.3.3.1.2 Tagged meshes

Before describing the set of subdivision rules, the description of the tagged meshes is described, which the algorithms accept as input. These meshes represent piecewise-smooth surfaces with features described above.

The complete list of tags is as follows. Edges can be tagged as *crease edges*. A vertex with incident crease edges receives one of the following tags:

- *Crease vertex*: joins exactly two incident crease edges smoothly.
- *Corner vertex*: connects two or more creases in a corner (convex or concave).

- *Dart vertex*: causes the crease to blend smoothly into the surface; this is the only allowed type of crease termination at an interior non-corner vertex.




Figure 5 — Features: (a) concave corner, (b) convex corner, (c) smooth boundary/crease, (d) corner with two convex sectors

Boundary edges are automatically assigned crease tags; boundary vertices that do not have a corner tag are assumed to be smooth crease vertices.

Crease edges divide the mesh into separate patches, several of which can meet in a corner vertex. At a corner vertex, the creases meeting at that vertex separate the ring of triangles around the vertex into sectors. Each sector of the mesh is labeled as *convex sector* or *concave sector* indicating how the surface should approach the corner.

In all formulas k denotes the number of faces incident at a vertex in a given sector. Note that this is different from the standard definition of the vertex degree: faces, rather than edges are counted. Both quantities coincide for interior vertices with no incident crease edges. This number is referred as *crease degree*.

The only restriction on sector tags is that concave sectors consist of at least two faces. An example of a tagged mesh is given in Figure 6.




Figure 6 — Crease edges meeting in a corner with two convex (light grey) and one concave (dark grey) sectors; the subdivision scheme modifies the rules for edges incident to crease vertices (e.g. e_1) and corners (e.g. e_2)

4.3.3.2 SubdivisionSurface

4.3.3.2.1 Node interface

SubdivisionSurface { #%NDT=SFGGeometryNode, SFBBaseMeshNode

eventIn	MFInt32	set_colorIndex
eventIn	MFInt32	set_coordIndex
eventIn	MFInt32	set_texCoordIndex
eventIn	MFInt32	set_creaseEdgeIndex
eventIn	MFInt32	set_cornerVertexIndex
eventIn	MFInt32	set_creaseVertexIndex
eventIn	MFInt32	set_dartVertexIndex

exposedField SFColorNode

color

NULL

exposedField	SFCoordinateNode	coord	NULL	
exposedField	SFTextureCoordinateNode	texCoord	NULL	
exposedField	SFInt32	subdivisionType	0	# 0 – midpoint, 1 – Loop, 2 – butterfly, # 3 – Catmull-Clark
exposedField	SFInt32	subdivisionSubType	0	# 0 – Loop, 1 – Warren-Loop
field	MFInt32	invisibleEdgeIndex	[]	# [0, ∞) # extended Loop if non-empty
exposedField	SFInt32	subdivisionLevel	0	# [-1, ∞)
# IndexedFaceSet fields				
field	SFBool	ccw	TRUE	
field	MFInt32	colorIndex	[]	# [-1, ∞)
field	SFBool	colorPerVertex	TRUE	
field	SFBool	convex	TRUE	
field	MFInt32	coordIndex	[]	# [-1, ∞)
field	SFBool	solid	TRUE	
field	MFInt32	texCoordIndex	[]	# [-1, ∞)
# tags				
field	MFInt32	creaseEdgeIndex	[]	# [-1, ∞)
field	MFInt32	dartVertexIndex	[]	# [-1, ∞)
field	MFInt32	creaseVertexIndex	[]	# [-1, ∞)
field	MFInt32	cornerVertexIndex	[]	# [-1, ∞)
# sector information				
exposedField	MFSubdivSurfaceSectorNode	sectors	[]	

4.3.3.2.2 Functionality and semantics

The **SubdivisionSurface** node is similar to the existing **IndexedFaceSet** node with all fields relating to face normals and the **creaseAngle** field removed since normals are generated by the subdivision algorithm at run-time and their behavior at facet boundaries is controlled by edge and vertex crease tagging. The control hull is specified as a collection of polygons and must be a manifold mesh. No “dangling” edges or vertices are permitted, and each polygon may share each of its edges with at most one other polygon. More details about the subdivision process and rules are described in subclause 4.3.3.2.3.

The **SubdivisionSurface** node represents a 3D shape formed through surface subdivision of a control hull constructed from vertices listed in the **coord** field. The **coord** field contains a **Coordinate** node that defines the 3D vertices referenced by the **coordIndex** field. **SubdivisionSurface** uses the indices in its **coordIndex** field to specify the polygonal faces of the control hull by indexing into the coordinates in the **Coordinate** node. An index of “-1” is used as a separator for the control hull faces. The last face may be (but does not have to be) followed by a “-1” index. If the greatest index in the **coordIndex** field is N, then the **Coordinate** node shall contain N+1 coordinates (indexed as 0 to N). Each face of the **SubdivisionSurface** control hull shall have:

- exactly three non-coincident vertices for triangular schemes (midpoint, Loop and butterfly);
- at least three non-coincident vertices for quadrilateral schemes (Catmull-Clark and extended Loop);
- vertices that define a planar, non-self-intersecting polygon.

Otherwise, the results are undefined.

The **SubdivisionSurface** node is specified in the local coordinate system and is affected by the transformation of its ancestors.

The **subdivisionType** field specifies the major subdivision scheme:

0. Midpoint (non-smoothing);
1. Loop (subclause 4.3.3.2.3.1);
2. Dyn's butterfly (cf. [36], [95]);
3. Catmull-Clark (subclause 4.3.3.2.3.3).

The **subdivisionSubType** field specifies a sub category of subdivision rules in the major subdivision rule. For this specification, for **subdivisionType** 0 (Loop), **subdivisionSubType** 0 indicates Loop and 1 Warren's Loop.

A **subdivisionType** value of zero and a non-empty **invisibleEdgeIndex** array imply that the extended Loop scheme (subclause 4.3.3.2.3.2) is used.

The **invisibleEdgeIndex** field is used to select co-ordinate pairs from **coord** to define which edges should be tagged as invisible. Consecutive pairs of co-ordinates define an edge; this pair of vertices must belong to one facet as defined by the **coordIndex** field and may be "interior" to the facet (e.g. in the case of a quadrilateral, the invisible edge would be one of its diagonals). All non-triangular facets should have enough tags to fully describe their triangulation. Note that a non-empty **invisibleEdgeIndex** array is only allowed when **subdivisionType** is 0; the results are undefined in other cases.

The **subdivisionLevel** field specifies the number of times the surface should be subdivided. If this field is "-1" then the terminal should apply enough iterations of subdivision until it reaches a "visually smooth" surface.

The fields **color**, **coord**, **texCoord**, **ccw**, **colorIndex**, **colorPerVertex**, **convex**, **coordIndex**, **solid**, and **texCoordIndex** have the same semantic as for the **IndexedFaceSet** node. These fields define the control hull of the subdivision surface. Textures coordinates and color information, if present, propagate from the control hull to subsequent subdivision levels through linear interpolation.

The **creaseEdgeIndex** field contains an indexed line set of creases, indicating where creases appear on the otherwise smooth subdivision surface (subclause 4.3.3.1.2). As for the **IndexedLineSet** node, the line is defined by successive vertex indices defined in the **coord** field. An index of "-1" indicates that the current crease has ended and the next one begins. The last crease may be (but does not have to be) followed by a "-1". Crease edges must already exist in the control hull topology (i.e. the edge must coincide with a single edge of a polygon as defined by the **coordIndex** field). Note that if an edge is tagged as both invisible and a crease, the results are undefined.

From the **creaseEdgeIndex** array, the default behavior (and appropriate subdivision rules) for vertices on crease edges are defined (subclause 4.3.3.1.2) as:

- Vertices with exactly one crease edge in their neighborhood are dart vertices,
- Vertices with exactly two crease edges in their neighborhood are crease vertices,
- Vertices with three or more crease edges in their neighborhood are corner vertices.

However, this array may not provide a description general enough for all surfaces and three arrays are provided: **cornerVertexIndex**, **creaseVertexIndex**, and **dartVertexIndex**. The information provided by these arrays completes (or overrides) the information in **creaseEdgeIndex**.

The **cornerVertexIndex** field indicates the vertex index of corner vertices.

The **creaseVertexIndex** field indicates the vertex index of crease vertices.

The **dartVertexIndex** field indicates the vertex index of dart vertices.

Finally, the **sectors** field may optionally contain an array of **SubdivSurfaceSector** nodes that describe additional sector information; see subclause 4.3.3.2.3 for more details. Note that if both the **sector** and the **invisibleEdgeIndex** fields are non-empty, the results are undefined.

4.3.3.2.3 Subdivision Stencils

In order to reduce the angles between adjacent triangles, the position of each new vertex introduced by the subdivision process is calculated as a function of some old vertices in the vicinity of the considered edge. The position of old vertices themselves may be altered as well for that same purpose.

The *stencil* of the subdivision scheme refers to the set of old vertices used to compute a new vertex. The term *mask*, or the set of (averaging) rules, is used to refer to the set of values of coefficients. It determines the convergence of the process and the smoothness degree of the limit surface.

Different masks are necessary for different types of vertices discussed in more detail in the next subclause.

4.3.3.2.3.1 Loop subdivision

Figure 7 shows the stencils used for repositioning old vertices and inserting new ones by splitting edges in the case of Loop's scheme [54].

Vertex rules are used to update positions of old vertices, edge rules are used to compute positions of newly inserted vertices.

For untagged interior and dart vertices rules shown in Figure 7a are used; for boundary/crease vertices rules in Figure 7b are used. Positions of corner vertices are never changed.




Figure 7 — Stencils of Loop's subdivision scheme: (a) repositioning interior old vertices (b) repositioning boundary or crease old vertices, (c) splitting interior edges not adjacent to boundary extraordinary vertices, (d) splitting boundary and crease edges, (e) splitting interior edges adjacent to an extraordinary (valence not 4) vertex on the boundary/crease or corner vertex; the values of beta and gamma are defined in the text

The coefficient β can be chosen in different ways.

Loop's original proposal was $\beta = \frac{1}{k} \left(\frac{5}{8} - \left(\frac{3}{8} + \frac{1}{4} \cos \frac{2\pi}{k} \right)^2 \right)$. In particular, this is $\beta = 3/16$ for $k = 3$, and $\beta = 1/16$ for

$k = 6$. Warren proposed later to keep $\beta = 3/16$ for $k = 3$, but to have simply $\beta = \frac{3}{8k}$ for $k > 3$, which also amounts to $\beta = 1/16$ for $k = 6$ [88].

Edge rules are more complex. The rule for an edge point inserted on an edge e is chosen depending on the tag of the edge and the tags of adjacent vertices and sectors. There are several cases:

1. In the absence of tags, the standard edge rules are applied (Figure 7c).
2. If a crease edge e is split, and no endpoint is a dart vertex, the rule shown in Figure 7d is used.
3. If one of the endpoints is a dart vertex, the standard rule (Figure 7c) is applied.
4. If one endpoint v is an extraordinary vertex (valence not 4) and tagged as corner or smooth crease and the edge e is not a crease edge, we use a modified rule shown in Figure 7e (with weight $3/4 - \gamma$ associated with extraordinary vertex), with parameter γ chosen as $\gamma(\theta_k) = 1/2 - 1/4 \cos \theta_k$. For a crease vertex, $\theta_k = \pi/k$, where k is the crease degree of the vertex in the sector where the edge e is.
5. If both endpoints are extraordinary vertices (valence not 4) and tagged as corner or smooth crease and the edge e is not a crease edge, then the average of the masks given by the rules in 4 above is taken.

At a corner vertex v , e is differentiated whether it is in a convex or concave sector. For a convex corner, $\theta_k = \alpha/k$, where α is the angle between the two crease edges spanning the sector, for concave corners, $\theta_k = (2\pi - \alpha)/k$.

For concave corners, one may use $\theta_k > \pi/k$ and for convex corners $\theta_k < \pi/k$. The simplest choice of θ_k is $3\pi/(2k)$ for concave corners and $\pi/(2k)$ for convex corners which is the default; the best choice is to use α/k for convex corners (respectively $(2\pi - \alpha)/k$ for concave corners) where α is the angle between edges of the control mesh bounding the sector; θ_k for concave and convex corners can be also a user-defined parameter changing in the ranges specified above.

Application of the stencils is followed by application of flatness modification described in the next subclause. This modification is optional for all cases excluding concave corner; in which case it is necessary to ensure tangent plane continuity of the surface.

(i) Flatness modification

With k defined as in previous subclause, let the vector of control points be $p^m = \{p_c^m, p_0^m, \dots, p_{k-1}^m\}$ (Figure 8). The superscript indicates the subdivision level the subscript indicates the point number in the sector. After subdivision step in a neighborhood of a corner vertex is modified as follows:

$$p^{new} = (1-s)p + s(\mathbf{a}_0x^0 + \mathbf{a}_1x^1 + \mathbf{a}_2x^2)$$

where s is the flatness parameter. $\mathbf{a}_i = (l^i, p)$, l^0, l^1, l^2 are limit position and tangent masks defined in the next subclause, and x^0, x^1, x^2 are the right eigenvectors also defined in the next subclause. The valid range of the flatness parameter is $0 \leq s \leq 1$ for all vertex types except concave corner, in which case it is $1 - \frac{1}{2\lambda} < s \leq 1$, with λ defined as $\lambda = 1/2 - 1/4(\cos\theta_k - \cos(\pi/k))$. Geometrically, the modified rule blends between control point positions before flatness modification and certain points in the tangent plane, which are typically close to the projection of the original control point. The limit position $\mathbf{a}_0$ of the center vertex remains unchanged.




Figure 8 — Neighborhoods of a vertex on different subdivision levels; the subdivision matrix relates the vector of control points p^m to the control points on the next level p^{m+1} ; for a neighborhood of k triangles $p^m = \{p_c^m, p_0^m, \dots, p_{k-1}^m\}$

The flatness modification is always applied at concave corner vertices; the default value for the flatness parameter is 1. Better visual results are obtained for $s = 1 - 1/(4\lambda)$. In other cases, s can be taken to be 0 by default and the modification step need not be applied.

Optionally, a user-specified flatness value can be associated with any vertex and used to control how fast the surface approaches the tangent plane.

A compliant implementation is required to support $s=0$ and $s=1$, but an implementation may support arbitrary s values.

(ii) Limit position and tangent masks

For each type of vertex (interior, smooth crease, corner, dart) there are masks for computing the limit position of this vertex and two tangent vectors at this vertex. The limit position masks are denoted l^0 and limit tangent position masks are denoted l^1 and l^2 respectively. A subscript c denotes the coefficient corresponding to the center vertex. Once two tangent vectors are obtained by applying the masks, the normal can be computed as the cross-product of the two. In addition we define right eigenvectors in each case which are necessary for flatness modification. Recall that dominant right eigenvector x^0 is the vector consisting of ones

Untagged interior or dart vertex of degree k . In all cases i is in the range $0 \dots k-1$,

and $\theta_k = 2\pi/k$.

$$l_c^0 = \frac{1}{1 + (8k/3)\beta}, \quad l_i^0 = \frac{(8/3)\beta}{1 + (8k/3)\beta}$$

$$x_c^1 = x_c^2 = l_c^1 = l_c^2 = 0$$

$$x_i^1 = \sin i\theta_k, \quad x_i^2 = \cos i\theta_k,$$

$$l_i^1 = \frac{2}{k} \sin i\theta_k, \quad l_i^2 = \frac{2}{k} \cos i\theta_k,$$

Smooth crease vertex of crease degree k . Let $\theta_k = \pi/k$

$$l_c^0 = 2/3, \quad l_1^0 = l_k^0 = 1/6, \quad l_i^0 = 0, \quad i = 1 \dots k-1$$

For $k=1$,

$$x_c^1 = -1/3, \quad x_1^1 = x_2^1 = 2/3, \quad x_c^2 = 0, \quad x_1^2 = 1, \quad x_2^2 = -1$$

$$l_c^1 = -1, \quad l_1^1 = l_2^1 = 1/2, \quad l_c^2 = 0, \quad l_1^2 = 1/2, \quad l_2^2 = -1$$

Otherwise $x_c^1 = x_c^2 = l_c^1 = 0$,

$$x_i^1 = \sin i\theta_k, \quad x_i^2 = \cos i\theta_k, \quad i = 0 \dots k$$

$$l_0^1 = 1/2, \quad l_k^1 = -1/2, \quad l_i^1 = 0, \quad i = 1 \dots k$$

$$l_c^2 = -\frac{2}{k} \left(\left(\frac{2}{3} - a \right) \sigma_1 - b \sigma_3 \right)$$

$$l_0^2 = l_k^2 = -\frac{2}{k} \left(\left(\frac{a}{2} + \frac{1}{6} \right) \sigma_1 + \frac{1}{2} b \sigma_3 \right)$$

$$l_i^2 = \frac{2}{k} \sin i\theta_k, \quad i = 1 \dots k-1$$

where

$$a = \frac{\frac{1}{4}(1 + \cos \theta_k)}{3(\frac{1}{2} - \frac{1}{4} \cos \theta_k)}, \quad b = \frac{\frac{2}{3} - a}{\cos \frac{k\zeta}{2}}$$

$$\sigma_1 = \frac{\sin \theta_k}{1 - \cos \theta_k}, \quad \sigma_3 = \frac{\cos \frac{k\zeta}{2} \sin \theta_k}{\cos \zeta - \cos \theta_k}$$

$$\zeta = \arccos(\cos \theta_k - 1)$$

Convex/concave corner vertex of crease degree k with parameter θ_k . Let $\theta = k\theta_k$.

$$l_c^0 = 1, \quad l_i^0 = 0, \quad i = 0 \dots k$$

$$x_c^1 = x_c^2 = 0, \quad x_i^1 = \frac{\sin i\theta_k}{\sin \theta}, \quad x_i^2 = \frac{\sin(k-i)\theta_k}{\sin \theta}, \quad i = 0 \dots k$$

$$l_c^1 = -1, \quad l_k^1 = 1, \quad l_i^1 = 0, \quad i = 0 \dots k-1$$

$$l_c^2 = -1, \quad l_0^2 = 1, \quad l_i^2 = 0, \quad i = 1 \dots k$$

4.3.3.2.3.2 Extended Loop subdivision

Extended Loop is an enhancement of the well-known triangle-based Loop subdivision algorithm that allows the designer to work with quadrilateral-based models which are better suited to capturing the symmetries of natural and man-made objects, while ensuring that the visualisation process can be optimised for a triangle-only rendering pipeline. Furthermore, unlike quadrilateral-based subdivision schemes such as Catmull-Clark, extended Loop offers the model designer exact control over the final model triangulation, which is essential for efficient rendering of low polygon models.

(i) Stencils

Extending Loop to handle arbitrary meshes requires the addition of edge visibility tags for edges that are introduced during triangulation and are not part of the original control hull. The introduction of edge visibility information will clearly affect the vertex and edge subdivision rules.




Figure 9 — Stencils of extended Loop subdivision for repositioning interior(smooth and dart) (a), (b) and boundary/crease old vertices and splitting interior (d), (e) and boundary/crease edges

Figure 9 shows the extended Loop subdivision stencils used for repositioning old vertices (as with standard Loop, corner vertices are never moved) and splitting edges both on the surface interior and boundaries/creases. Solid lines indicate original control hull (visible) edges while dashed lines indicate invisible edges. It is clear from the stencils (a), (b) and (c) that repositioning old vertices with extended Loop is identical to standard Loop if invisible edges are ignored, with one exception (b). If a vertex has two or less visible edges in its neighborhood, then the invisible edges are included in the calculation. If all edges are visible then the stencils are identical to standard Loop subdivision.

As with standard Loop subdivision, $\beta = \frac{1}{k} \left(\frac{5}{8} - \left(\frac{3}{8} + \frac{1}{4} \cos \frac{2\pi}{k} \right)^2 \right)$, where k is the number of visible edges (Figure

9a) or the valence (visible + invisible edges) if the number of visible edges is ≤ 2 (Figure 9b).

The stencil for splitting boundary/crease edges (Figure 9f) is also identical to standard Loop, while splitting interior edges depends on whether the edge in question is visible (Figure 9d – left) or not (Figure 9d – right). As with the old vertex repositioning stencils, the edge split stencils try to employ vertices at the “other end” of visible edges only, ignoring invisible edges completely. Identifying visible edges at each side of the edge endpoint involves a clockwise and/or counter-clockwise search around the endpoint neighborhood and may result in the same vertex for both sides or no visible edges in which case the endpoint is assigned the full weight (Figure 9e). Again if all edges are visible, the stencils are identical to standard Loop subdivision.

(ii) Visibility propagation

The visibility of the new edges needs to be determined, since it will be required on the next round of subdivision. New edges radiate from a new point (see Figure 10), and their visibility is determined according to the neighborhood of the point.




Figure 10 — New edges created when splitting an edge in two

Two edges will be the two halves of the original edge E - these simply keep the visibility of the original edge. Of the other four new edges, two edges will be to the right hand side of the original edge, and two more to its left. If we label these edges clockwise as e_1 , e_2 , e_3 and e_4 , these connect to other new points created by dividing the original edges labeled E_1 , E_2 , E_3 and E_4 . The visibility of the new edges is initialized to invisible and is changed to visible if the lookup of the combination of E , E_1 , E_2 , E_3 and E_4 in Table 1 indicates that it should be. Clearly, the visibility of each of e_1 , e_2 , e_3 and e_4 is not uniquely determined by splitting E alone; e_1 , for example, also depends on the visibility lookup for splitting E_1 .

Table 1 — New edge visibility lookup table

Original edges visible					New edges visible			
E	E_1	E_2	E_3	E_4	e_1	e_2	e_3	e_4
				✓	✓		✓	
			✓			✓		✓
			✓	✓	✓	✓		
		✓			✓		✓	
		✓		✓	✓		✓	
		✓	✓		✓			✓
		✓	✓	✓	✓			
	✓					✓		✓
	✓			✓		✓	✓	
	✓		✓			✓		✓
	✓		✓	✓		✓		
	✓	✓					✓	✓
	✓	✓		✓			✓	
	✓	✓	✓				✓	✓
	✓	✓	✓	✓	✓	✓	✓	✓

Original edges visible					New edges visible			
E	E_1	E_2	E_3	E_4	e_1	e_2	e_3	e_4
✓								
✓				✓				
✓			✓					
✓			✓	✓			✓	✓
✓		✓						
✓		✓		✓				
✓		✓	✓					
✓		✓	✓	✓			✓	✓
✓	✓							
✓	✓			✓				
✓	✓		✓					
✓	✓		✓	✓			✓	✓
✓	✓	✓			✓	✓		
✓	✓	✓		✓	✓	✓		
✓	✓	✓	✓		✓	✓		
✓	✓	✓	✓	✓	✓	✓	✓	✓

Note that Table 1 is unchanged if we swap left and right, so it is unimportant which side of the original edge we call "right" to start with.

(iii) Vertex normal calculation

In order to perform operations such as lighting on a surface, it is necessary to determine the normal to that surface at each vertex. This is the direction at right angles to the limit surface at that point. In order to calculate this normal, it is necessary first to obtain two tangent vectors, t_1 and t_2 , which are directions that lie along the surface rather than at right angles to it. Given these, a vector cross product will result in the desired normal:

$$n = t_1 \times t_2.$$




Figure 11 — (a) interior, and (b) boundary/crease vertex neighborhood used for normal calculation

For interior smooth and dart vertices of valence k (Figure 11a), the tangents are defined as the weighted averages of the positions of the neighborhood vertices (v_i where $i \in [1, k]$):

$$\begin{aligned} t_1 &= c_1 v_1 + c_2 v_2 + c_3 v_3 + \dots + c_k v_k \\ t_2 &= c_2 v_1 + c_3 v_2 + c_4 v_3 + \dots + c_k v_{k-1} + c_1 v_k, \end{aligned}$$

where:

$$c_i = \cos \frac{2\pi i}{k}.$$

For boundary, crease and corner vertices (Figure 11b), each sector between pairs of boundary/crease edges will have a different pair of tangent vectors that are defined as:

$$t_1 = v_1 - v_k$$

$$t_2 = \begin{cases} -2v_0 + v_1 + v_2, & k = 2 \\ v_2 - v_0, & k = 3, \\ -2v_0 - v_1 + 2v_2 + 2v_3 - v_4, & k = 4 \\ w_1 v_1 + w_2 v_2 + w_3 v_3 + \dots + w_k v_k, & k > 4 \end{cases}$$

where, for $\theta = \frac{\pi}{k-1}$, w_i is defined as:

$$w_i = \begin{cases} \sin \theta, & i = 1, i = k \\ (2 \cos \theta - 2) \sin((i-1)\theta), & 2 \leq i \leq k-1 \end{cases}$$

4.3.3.2.3.3 Catmull-Clark subdivision

The Catmull-Clark scheme was described in [23]. This scheme can be applied to general polygonal meshes, but after the first subdivision step all polygons in a mesh become quads. The masks for a mesh consisting only of quads are shown in Figure 12 and Figure 13. The scheme produces surfaces that are C^2 everywhere except at extraordinary vertices, where they are C^1 .




Figure 12 — Catmull-Clark rules with k defined as in subclause 4.3.3.1.2 (a) face rule; (b) regular edge rule

(c) boundary/crease edge rule (d) interior vertex rule $\beta = \frac{3}{2k}$ and $\delta = \frac{1}{4k}$ (e) smooth boundary/crease vertex rule

There are three types of rules: vertex rules used to update old vertex positions, edge rules used to insert new vertices for each edge, and face rules used to insert new vertices for each face.




Figure 13 — Modified rule for odd vertices adjacent to a boundary extraordinary vertex (Catmull-Clark scheme), the coefficient gamma is defined in the text

The rules used only on the first step are:

— a face control point for an n -gon is computed as the average of the corners of the polygon;

- an edge control point for an untagged edge as the average of the endpoints of the edge and newly computed face control points of adjacent faces; for a crease/boundary edge it is the average of the endpoints;
- The vertex rule for untagged interior and dart vertices is:

$$p_c^{j+1} = \frac{k-2}{k} p_c^j + \frac{1}{k^2} \left(\sum_{i=0}^{k-1} p_i^j + q_i^{j+1} \right)$$

where p_i^j are k control points at level j sharing edges with p_c^j , and q_i^{j+1} are face control points for faces incident at the vertex. Note that face control points on level $j+1$ are used.

- For boundary/crease vertices, it is the spline rule identical to the one used for the Loop scheme. The corner vertex positions are never changed.

After the first subdivision step the mesh consists only of quadrilaterals and the subdivision rules are described in (Figure 12 and Figure 13). The vertex and face rules are essentially the same; the masks are shown in Figure 12. As in the case of the Loop scheme edge rules are more complex:

- In the absence of tags, the standard edge rules are applied (Figure 14b).
- If a crease edge e is split and no endpoint is tagged as a dart vertex, rules of Figure 14c are used.
- If one of the endpoints is a dart vertex, the standard edge rule is used (Figure 14b).
- If one endpoint v is tagged as a corner or smooth crease and the edge e is not a crease, a modified rule shown in Figure 6 is used, with parameter γ chosen as $\gamma(\theta_k) = 3/8 - 1/4 \cos \theta_k$. Exactly like for Loop scheme, for a smooth crease vertex, $\theta_k = \pi/k$, where k is the crease degree of v in the sector where e is. For a corner use $\theta_k = \alpha/k$, where $\alpha > \pi$ for concave corners and $\alpha < \pi$ for convex corners.
- If both endpoints are tagged as corner or smooth crease, then the average of the masks given by the rules above is taken.

Application of the stencils is followed by application of flatness modification described in the next subclause. This modification is optional for all cases excluding concave corner; in this case it is necessary to ensure tangent plane continuity of the surface.

(i) Flatness modification

The flatness modification is similar to the one used for Loop scheme, but the vector of control points has to include one more group of points $[q_0^m, \dots, q_{k-1}^m]$ as shown in Figure 14. The superscript indicates the subdivision level the subscript indicates the point number in the sector.




Figure 14 — Neighborhoods of a vertex on different subdivision levels; the subdivision matrix relates the vector of control points p^m to the control points on the next level p^{m+1} ; for a neighborhood of k quadrilaterals $p^m = \{p_c^m, p_0^m, \dots, p_{k-1}^m, q_0^m, \dots, q_{k-1}^m\}$

The rule itself coincides with the rule for the Loop scheme:

$$p^{new} = (1-s)p + s(a_0x^0 + a_1x^1 + a_2x^2)$$

where s is the flatness parameter. $a_i = (l^i, p)$, l^0, l^1, l^2 are limit position and tangent masks defined in the next subclause, and x^0, x^1, x^2 are the right eigenvectors also defined in the next subclause..

The valid range of the flatness parameter is $0 \leq s \leq 1$ for all vertex types except concave corner, in which case it is $1 - \frac{1}{2\lambda} < s \leq 1$, with λ defined as

$$\lambda = \frac{1}{2}\gamma + \frac{1}{8}c^2 + \frac{1}{16}\sqrt{64\gamma^2 + 32c^2\gamma - 48\gamma + 4c^4 + 4c^2 + 9} \text{ and } c = \cos \frac{\pi}{2k}.$$

The default flatness parameter can be set to zero for all cases except concave corners, in which case it is set to 1. Better visual results are obtained for the values close to $s = 1 - 1/(4\lambda)$.

Optionally, a user-specified flatness value can be associated with any vertex and used to control how fast the surface approaches the tangent plane.

(ii) Limit position and tangent masks

Interior or dart vertex of degree k . In all cases i is in the range $0 \dots k-1$, and $\theta_k = 2\pi/k$.

$$\begin{aligned}
l_c^0 &= \frac{k}{k+5}, \quad l_{pi}^0 = \frac{4}{k(k+5)}, \quad l_{qi}^0 = \frac{1}{k(k+5)} \\
x_c^1 &= x_c^2 = l_c^1 = l_c^2 = 0 \\
x_{pi}^1 &= \frac{1}{\sigma} \sin i \theta_k, \quad x_{pi}^2 = \frac{1}{\sigma} \cos i \theta_k \\
x_{qi}^1 &= \frac{\sin i \theta_k + \sin(i+1) \theta_k}{\sigma(4\lambda-1)}, \quad x_{qi}^2 = \frac{\cos i \theta_k + \cos(i+1) \theta_k}{\sigma(4\lambda-1)}, \\
l_{pi}^1 &= 4 \sin i \theta_k, \quad l_{pi}^2 = 4 \cos i \theta_k, \\
l_{qi}^1 &= \frac{\sin i \theta_k + \sin(i+1) \theta_k}{4\lambda-1}, \quad l_{qi}^2 = \frac{\cos i \theta_k + \cos(i+1) \theta_k}{4\lambda-1}
\end{aligned}$$

where $\lambda = 5/16 + 1/16(\cos \theta_k + \cos \theta_k / 2\sqrt{9 + \cos 2\theta_k})$ and $\sigma = k(2 + \frac{1 + \cos \theta_k}{(4\lambda-1)^2})$

Smooth crease vertex of crease degree k . Let $\theta_k = \pi/k$

$$\begin{aligned}
l_c^0 &= 2/3, \quad l_{p1}^0 = l_{pk}^0 = 1/6, \\
l_{q0}^0 &= l_{qi}^0 = l_{pi}^0 = 0, \quad i = 1 \dots k-1
\end{aligned}$$

For $k = 1$,

$$\begin{aligned}
x_c^1 &= 1/18, \quad x_{p0}^1 = -2/18, \quad x_{p1}^1 = -2/18, \quad x_{q0}^1 = -5/18 \\
x_c^2 &= 0, \quad x_{p0}^2 = -1/2, \quad x_{p1}^2 = 1/2, \quad x_{q0}^2 = 0 \\
l_c^1 &= 6, \quad l_{p0}^1 = -3, \quad l_{p1}^1 = -3, \quad l_{q0}^1 = 0 \\
l_c^2 &= 0, \quad l_{p0}^2 = -1, \quad l_{p1}^2 = 1, \quad l_{q0}^2 = 0
\end{aligned}$$

Otherwise $l_c^2 = x_c^1 = x_c^2 = 0$,

$$\begin{aligned}
x_{pi}^1 &= \sin i \theta_k, \quad x_{pi}^2 = \cos i \theta_k, \quad i = 0 \dots k \\
x_{qi}^1 &= \sin i \theta_k + \sin(i+1) \theta_k, \quad i = 0 \dots k-1 \\
x_{qi}^2 &= \cos i \theta_k + \cos(i+1) \theta_k, \quad i = 0 \dots k-1 \\
l_{p0}^2 &= 1/2, \quad l_{pk}^2 = -1/2, \quad l_{q0}^2 = l_{qi}^2 = l_{pi}^2 = 0, \quad i = 1 \dots k-1 \\
l_c^1 &= 4R(\cos \theta_k - 1), \quad l_{p0}^1 = l_{pk}^1 = -R(1 + 2 \cos \theta_k) \\
l_{pi}^1 &= \frac{4 \sin i \theta_k}{(3 + \cos \theta_k)k}, \quad i = 1 \dots k-1 \\
l_{qi}^1 &= \frac{4(\sin i \theta_k + \sin(i+1) \theta_k)}{(3 + \cos \theta_k)k}, \quad i = 0 \dots k-1
\end{aligned}$$

where $R = \frac{\cos \theta_k + 1}{k \sin \theta_k (3 + \cos \theta_k)}$

Convex/concave corner vertex of crease degree k with parameter θ_k . Let $\theta = k\theta_k$.

Left eigenvectors are the same as for Loop with zeroes everywhere except l_c , l_{p0} and l_{pk} .

$$\begin{aligned}
x_{pi}^1 &= \frac{\sin i \theta_k}{\sin \theta}, \quad x_{pi}^2 = \frac{\sin(k-i) \theta_k}{\sin \theta}, \quad i = 0 \dots k \\
x_{qi}^1 &= \frac{\sin i \theta_k + \sin(i+1) \theta_k}{\sin \theta}, \quad i = 0 \dots k-1 \\
x_{qi}^2 &= \frac{\sin(k-i) \theta_k + \sin(k-i-1) \theta_k}{\sin \theta}, \quad i = 0 \dots k-1
\end{aligned}$$

4.3.3.2.3.4 Normal control

In addition to the flatness modification which insures that the surface is tangent plane continuous at concave corner vertices and provides additional shape control at other vertex types, a similar modification mechanism can be used to control the tangent plane of the surface as shown in Figure 15.

Specifically, if the user prescribes tangent vectors or the normal for a sector at a corner or smooth boundary vertex or a normal for an interior vertex then the following modification can be applied after each subdivision step for both Loop and Catmull-Clark surfaces,

$$p^{new} = p + t((\mathbf{a}'_1 - \mathbf{a}_1)x^1 + (\mathbf{a}'_2 - \mathbf{a}_2)x^2)$$

where the notation used is the same as in the subclauses describing flatness modification above, $\mathbf{a}'_1$ and $\mathbf{a}'_2$ denote the prescribed tangents and t is a user-specified parameter $0 \leq t \leq 1$, by default set to 1. If only the normal is prescribed, the tangents are computed as $\mathbf{a}'_i = \mathbf{a}_i - (\mathbf{a}_i, \mathbf{n})\mathbf{n}$ $i = 1, 2$.

The parameter t determines how fast the surface approaches the prescribed tangent plane.




Figure 15 — Normal interpolation: (a) Surface with convex corners; (b) Prescribed directions: at each corner we tilt the normal for one surface sector slightly inwards; (c) Smooth surface; (d) Same control mesh but all normals vertical

4.3.3.3 SubdivSurfaceSector

4.3.3.3.1 Node interface

```
SubdivSurfaceSector { #%NDT=SFSubdivSurfaceSectorNode
    field          SInt32    faceIndex    0 # face index from coordIndex
    field          SInt32    vertexIndex  0 # vertex index from coord
    exposedField   SInt32    tag          0
    exposedField   SFloat    flatness     0
    exposedField   SFloat    theta        0
    exposedField   SVec3f    normal       0 0 0
    exposedField   SFloat    normalTension 0
}
```

4.3.3.3.2 Functionality and semantics

The **faceIndex** field contains the index of the face and the **vertexIndex** field contains the vertex index the sector is attached to. The indices refer to point index and face index from **coord** and **coordIndex** fields respectively in the parent **SubdivisionSurface** node.

The **tag** field indicates the type of sector (Figure 16):

Tag value	Semantic
0	No corner sector
1	Convex sector
2	Concave sector




Figure 16 — Concave corners (left) and convex corners (right)

The **flatness** field contains a value between 0 and 1: 0 indicating no modified flatness, and 1 indicating maximal flatness. This parameter is denoted s in Loop and Catmull-Clark rule definitions above (Figure 17).




Figure 17 — Influence of the flatness parameter

The **theta** field is an angle used to control the shape of the corner; this angle is denoted θ_k in the mask definitions above (Figure 18).



Figure 18 — Influence of theta

The **normal** field specifies the desired normal at the vertex (Figure 19).

The **normalTension** field is the influence of the normal between 0 and 1, 0 meaning no influence; denoted t in the formulas for normal modification (Figure 19).




Figure 19 — Prescribed normals and tension parameter

4.3.3.4 WaveletSubdivisionSurface

4.3.3.4.1 Node interface

WaveletSubdivisionSurface { #NDT=%SF3DNode

exposedField	SFGeometryNode	baseMesh	NULL	
exposedField	SFFloat	frequency	1.0	
exposedField	SFFloat	fieldOfView	0.785398	# pi/4
exposedField	SFInt32	quality	1	

}

4.3.3.4.2 Functionality and semantics

Wavelet representation of surfaces relies on an extension of classical multiresolution analysis theory in which operators similar to those of subdivision surfaces are used in place of traditional interval subdivision and low-pass filtering (scaling spaces).

In this representation, a 3D mesh is received as a coarse mesh and independent sets of zero-tree encoded wavelet coefficients, allowing adaptive and progressive reconstruction at the client side.

The **WaveletSubdivisionSurface** node is used to represent adaptively multi-resolution polygonal models. It allows representing at any moment a 3D-model at a resolution that depends on the distance of the object with respect to the viewpoint. Transmission of such a model consists of two parts: a base mesh, which describes the model at the coarsest level of detail, and series of refinements (wavelet coefficients) that can be carried on one or more elementary streams (see the bitstream definition in subclause 5.3.1).

The **baseMesh** field shall specify either a **SubdivisionSurface** node or an **IndexedFaceSet**. If **baseMesh** is a **SubdivisionSurface**, the stream containing the mesh shall consist of a single access unit, and it will be explicitly instantiated as a **SubdivisionSurface**. The structure of this **SubdivisionSurface** will be preserved, namely deletion and insertion of new faces and/or vertices is explicitly forbidden. However, replacing the values of the coordinates is allowed. If **baseMesh** is an **IndexedFaceSet**, it is required to be manifold.

NOTE — Instantiating explicitly **baseMesh** as a **SubdivisionSurface** or **IndexedFaceSet** allows animating the node.

The terminal is responsible for maintaining an internal representation of the mesh to be rendered consisting of the **baseMesh** to which the **downstream** data must be applied.

The **frequency** field gives a preferred frequency at which the terminal shall give its state by the means of the back channel stream if it exists.

The **quality** field allows parameterizing the quality of the rendered object. The range of values for this field is [0...2] where 0 stands for low quality, 1 for normal quality, 2 for best quality. Those values don't have a more explicit semantic, that is, it is a hint for the terminal.

The player sees the **WaveletSubdivisionSurface** node as a **SubdivisionSurface** or **IndexedFaceSet** node container. That is, the player has access to the **baseMesh** field.

4.3.4 MeshGrid representation

4.3.4.1 Introduction

The peculiarity of the MeshGrid representation [73] is that it defines the object's wireframe (see e.g. Figure 20(a)) by (1) describing the connectivity between the vertices located on the surface of the object by a so-called connectivity-wireframe (see Figure 20(b)), and (2) positioning these vertices in relation to a regular 3-D grid of points, i.e. the reference-grid (see Figure 20(c)).

Both the connectivity-wireframe and the reference-grid can be represented hierarchically, as shown in Figure 20(d) and Figure 20(e) respectively for 3 levels in the hierarchy.




Figure 20 — A MeshGrid humanoid model designed by means of implicit surfaces

4.3.4.1.1 The reference-grid (RG) – the reference-system of the MeshGrid model

The reference-grid (see Figure 20(c)) is the reference system upon which the MeshGrid model is built, and is defined by the intersection points between three sets of reference-surfaces S_U, S_V, S_W as given by the formula (1)

$$RG = \bigcap \left\{ \sum_U S_U, \sum_V S_V, \sum_W S_W \right\}. \quad (1)$$

The discrete position (u, v, w) of a reference-grid point represents the indices of the reference-surfaces $\{S_U, S_V, S_W\}$ intersecting in that point, while the coordinate (x, y, z) of a reference-grid point is equal to the coordinate of the computed intersection point.

A multi-resolution model is designed by choosing a hierarchical reference-system, the reference-surfaces belonging to the same resolution level being displayed in the same color in the example from Figure 20(e).

The reference system should be chosen according to the topology of the object, i.e. it has higher density in (1) areas where the curvature is high, or (2) places that are deformed during the animation. For the humanoid model of Figure 20, the reference system is adjusted for traversing the anatomical articulations (joints) of the body, in order to be able to virtually split the resulting seamless connectivity-wireframe into meaningful anatomical parts (such as the shoulder, elbow, wrist).

4.3.4.1.2 The connectivity-wireframe (CW) – the surface mesh

The connectivity-wireframe (see Figure 20(b)) is generated by merging the contouring of the original 3D model in each of the reference-surfaces S_U, S_V, S_W as given by the formula (2)

$$CW = \bigcup \left\{ \sum_U C(S_U), \sum_V C(S_V), \sum_W C(S_W) \right\}, \quad (2)$$

in which $C(S_x)$ represents the contour obtained by the intersection with surface S_x . Each vertex is located at the intersection point between two contours.

Subdividing the reference surfaces hierarchically, as explained in subclause 4.3.4.1.1, results in a multi-resolution connectivity-wireframe, in which each resolution level consists of a single seamless mesh, as shown in Figure 20(d).

4.3.4.1.3 The features

This type of representation is specific for describing objects defined as series of contours or slices, such as: discrete 3D data sets (e.g. 3D medical images, processed range scanner data), cylindrical and spherical projected scanner meshes based on quadrilaterals, or generic models (see Figure 21(a)-(d)).




Figure 21 — The rendered surface obtained from (a) a 3D voxel model of the brain, (b) a structured light scanner quadrilateral mesh, respectively with its deformed reference-grid in overlay in (c), and (d) a generic model

A MeshGrid mesh allows view-dependent streaming (see subclause 5.1) and particular animation possibilities, such as: hierarchical grid animation, and offset based animation (see subclause 4.3.4.3).

4.3.4.2 MeshGrid

4.3.4.2.1 Node specification

MeshGrid {			
eventIn	MFInt32	set_colorIndex	
eventIn	MFInt32	set_coordIndex	
eventIn	MFInt32	set_normalIndex	
eventIn	MFInt32	set_texCoordIndex	
exposedField	SFColorNode	color	NULL
exposedField	SFCoordinateNode	coord	NULL
exposedField	SFInt32	displayLevel	0
exposedField	SFInt32	filterType	0
exposedField	SFCoordinateNode	gridCoord	NULL
exposedField	SFInt32	hierarchicalLevel	0
exposedField	MFInt32	nLevels	[]
exposedField	SFNormalNode	normal	NULL
exposedField	MFInt32	nSlices	[]
exposedField	SFTextureCoordinateNode	texCoord	NULL
exposedField	MFInt32	vertexLink	[]
exposedField	MFFloat	vertexOffset	[]
field	MFInt32	colorIndex	[]
field	MFInt32	coordIndex	[]
field	MFInt32	normalIndex	[]
field	SFBool	solid	TRUE
field	MFInt32	texCoordIndex	[]
eventOut	SFBool	isLoading	
eventOut	MFInt32	nVertices	
}			

4.3.4.2.2 Node semantics

The **MeshGrid** node is derived from the existing **IndexedFaceSet** node. The following fields have been removed, since the triangulation of the surface is pre-defined in the decoder: *ccw*, *convex*, *creaseAngle*. Some new fields are introduced, specific to the MeshGrid representation. The mesh is manifold, but can be open or closed.

The fields of the MeshGrid node can be encoded either (1) by BIFS, or (2) as a binary stream. When encoded as a binary stream, the following fields of the MeshGrid node will be initialized when decoding the binary stream (see subclause 5.2.3): **nSlices**, **nLevels**, **gridCoord**, **coord**, **coordIndex**, **vertexOffset**, **vertexLink**, **nVertices**. The **isLoading** flag is set to true while decoding the binary stream, respectively becomes false when the decoding finishes. The value of **isLoading** flag has to be checked before updating the fields in order to make sure that the decoder has finished modifying them. If encoded by BIFS only the **vertexLink** (the description of the connectivity-wireframe) and **vertexOffset** (the vertices' offsets relative to the reference-grid) fields need to be decoded, as explained in subclauses 5.2.3.1 and 5.2.3.3.

The **nSlices** field specifies the number of slices of the three sets $\{S_U, S_V, S_W\}$ of reference-surfaces defining the reference-grid at the last resolution level. The number of slices in the *U* direction equals to **nSlices[0]**, respectively the number of slices in the *V* direction equals to **nSlices**, and the number of slices in the *W* direction equals to **nSlices**. The minimum number of slices in any direction is equal to 1. More detail can be found in B.1.1.

The **nLevels** field defines the number of resolution levels of the MESHGRID mesh. The number of resolution levels in the *U* direction equals to **nLevels[0]**, respectively the number of resolution levels in the *V* direction equals to **nLevels**, and the number of resolution levels in the *W* direction equals to **nLevels**. More detail can be found in B.1.1.

The **gridCoord** field defines the reference-grid points. The number of reference-grid points can be either equal (1) to '8', or (2) to the following value: $nSlices[0] \times nSlices \times nSlices$. In case (1), the values represent the coordinates of the '8' corners of the reference-grid, and the decoder will generate a regular distributed reference-grid based on the position of the '8' corners. In this case the **gridCoord** buffer is resized and its initial values are replaced by the computed coordinates of the grid points. In case (2) the values represent the coordinates of the reference-grid points. Any point from the reference-grid can be addressed by a distinct (u, v, w) discrete position, where $\{u, v, w\}$ are integer values in the range: $u \in \{0, \dots, nSlices - 1\}$, $v \in \{0, \dots, nSlices - 1\}$, and $w \in \{0, \dots, nSlices - 1\}$. Given *u* the row index, *v* the column index, and *w* the plane index, the ordering of the reference-grid points in the **gridCoord** buffer is row first, column second, and plane third. When the MESHGRID node is initialized from a binary stream, then the contents of **gridCoord** is ignored, and it will be overridden with the coordinates decoded from the binary stream as explained in subclause 5.2.3.2. When the **gridCoord** field is updated by BIFS, the coordinates of the grid points that did not change can be updated in a hierarchical way as explained in subclause 4.3.4.3.1.

The **vertexLink** field defines the connectivities between the vertices in the connectivity-wireframe. The decoding of the contents of this field is explained in subclause 5.2.3.1. If the node is initialized from a binary stream, then the content of vertexLink is ignored.

The **vertexOffset** field specifies the scalar offset for each vertex relative to a corresponding reference-grid point. The meaning of the offset is explained in subclause 4.3.4.3.2, while the decoding of this field is described in subclause 5.2.3.3. The value of the scalar offset lies in the range $[0, 1)$. If the node is initialized from a binary stream, then the content of vertexOffset is ignored, and it will be overridden with the scalar offsets decoded from the binary stream.

The **coord** field is filled with the coordinates (x, y, z) of the vertices, as result of the decoding procedure (see subclause 5.2.3.1). Its contents, if any, will always be overridden. The coordinates of the vertices can be derived from the coordinates of the reference-grid points as explained in subclause 5.2.3.1.

The **coordIndex** field is filled with the indices of the triangles when building the triangulation of the decoded mesh as explained in 5.2.3.1.3. Its contents, if any, will always be overridden.

The **nVertices** field stores the number of vertices for each resolution level. This field is updated after the decoding procedure (see subclause 5.2.3.1).

The **hierarchicalLevel** field defines the resolution level of the reference-grid points (**gridCoord** field) that may change during the animation. The usage is explained in subclause 5.2.3.1.

The **filterType** field defines the type of filter used for the hierarchical interpolation of the reference-grid points during the animation.

filterType	Meaning
00	Use short filter (see subclause 5.2.3.1)
01	Use smooth filter (see subclause 5.2.3.1)
10	Reserved
11	Reserved

The **displayLevel** field specifies the resolution level of the MeshGrid model to be rendered. Accepted values are in the range $[0, \text{maxLevel}]$, where **maxLevel** is the maximum value between **nLevels[0]**, **nLevels**, **nLevels**. More detail can be found in Annex B.1.1.

The fields **normal**, **normalIndex**, **texCoord**, **texCoordIndex**, **color**, **colorIndex**, and **solid** have the same semantic as for the **IndexedFaceSet** node.

4.3.4.3 Animation Extensions

4.3.4.3.1 Hierarchical interpolation of the reference-grid points

In case of a multi-resolution MeshGrid model, animating the hierarchical reference-grid can be done on a hierarchical basis, more precisely any change in the coordinates of reference-grid points, i.e. the **gridCoord** field (see subclauses 4.3.4.2.1 and 4.3.4.2.2), at a certain resolution level will trigger a local update of the coordinates of the reference-grid points belonging to a higher resolution level. The new positions of the reference-grid points belonging to resolution level $(l+1)$ can be computed from the new positions of the neighboring reference-grid points at resolution level (l) via an interpolation method based on “Dyn’s four point scheme for curves” [35]. The position of a reference-grid point P_{2n}^{l+1} (shown in Figure 22) is computed as follows:

$$P_{2n}^{l+1} = \Delta P_{2n}^{l+1} + \left(-w \times P_{2n-3}^l + \left(\frac{1}{2} + w \right) \times P_{2n-1}^l + \left(\frac{1}{2} + w \right) \times P_{2n+1}^l - w \times P_{2n+3}^l \right), \quad (3)$$

$$\Delta P_{2n}^{l+1} = P_{2n}^{l+1} - \left(-w \times P_{2n-3}^l + \left(\frac{1}{2} + w \right) \times P_{2n-1}^l + \left(\frac{1}{2} + w \right) \times P_{2n+1}^l - w \times P_{2n+3}^l \right), \quad (4)$$

where l represents the hierarchical level of the grid point, P the last position of a point, P the new position, ΔP the detail computed for the last position P . The detail ΔP is added to the interpolated value in point P . The weight w defines the smoothness of the limit curve, and can be specified via the **filterType** field of the **MeshGrid** node. The weight w may have one of the following two values: $w = 0$ when **filterType** = 0 and $w = 1/16$ when **filterType** = 1. In general w is taken equal to $1/16$ which corresponds to fitting a Catmull-Rom or Cardinal spline curve through the points.




Figure 22 — Dyn’s four point scheme for curves applied for the hierarchical MeshGrid

The **hierarchicalLevel** field specifies the resolution level at which the grid points, i.e. the **gridCoord** field, may change. The interpolation of the grid points will take place only for the range of resolution levels [**hierarchicalLevel** + 1, **nLevels** - 1]. If the value of **hierarchicalLevel** $\geq$ **nLevels** - 1 than no interpolation occurs. In the third case, when **hierarchicalLevel** = -1, grid points may change at any resolution level simultaneously, the points that change are detected, and the interpolation of the remaining grid points is done according to equation (3).


The position of the vertices is updated from the grid positions as explained in subclause 4.3.4.3.2.

4.3.4.3.2 Relationship between the reference-grid points and vertices

The relationship between the vertices of the connectivity-wireframe and reference-grid points is explained in the 2D cross-section of Figure 23, in which the reference-grid is put on top of the object's section.

Any reference-grid line l (label 1) is the result of the intersection between two reference-surfaces S_1 and S_2 , belonging to two different sets (see subclause 4.3.4.1.1). Every vertex V , lying on a contour (label 2) of the object, stores the discrete position (u,v,w) – which has been derived during decoding – of the grid point it is attached to. The vertices are related to the reference-grid points in two different ways, as follows:

- 1) As shown in Figure 23(a)-(b), vertex V is positioned in-between two grid points, one inside the object G_1 to which it is attached, and one outside the object G_2 . The relative position of vertex V with respect to G_1 and G_2 is given by the scalar *offset* (label 3) expressed by equation (5). Because *offset* indicates the fractional position between G_1 and G_2 , any displacement of G_1 or G_2 (see the example of Figure 23 (b)) is automatically reflected in an update of the coordinates (x,y,z) of V through equation (6).
- 2) As illustrated in Figure 23(c), vertex V is attached in this case to a grid position G belonging to a *single-layer* reference-grid, i.e. the number of slices is equal to 1 in one of the $\{u,v,w\}$ directions. The relative position of vertex V with respect to G is given by the scalar *offset* (label 3) expressed by equation (7). Similarly to the observation made at the previous point, any displacement of G is automatically reflected in an update of the coordinates (x,y,z) of V through equation (8). In this case the displacements along the normal follow both directions as shown for the vertices with labels 4 and 5.



In this figure O denotes the origin of the reference coordinate system in which the reference-grid is defined.

Figure 23 — A cross-section through a 3D object, illustrating the contour of the object, the reference-grid, and the relation between the vertices (belonging to the connectivity-wireframe and located at the surface of the object) and the grid points

The offsets of the vertices are stored in the **vertexOffset** field of the **MeshGrid** node (see subclause 4.3.4.2).

4.3.4.4 Examples

The following example loads a VRML MeshGrid node.

```
Shape {
  appearance Appearance {
    material Material {
      diffuseColor 1.0 0.0 0.0
      ambientIntensity 0.1
    }
  }
  geometry MeshGrid {
    nSlices [ 3 3 3 ]
    nLevels 1
    displayLevel
    hierarchicalLevel 0
    gridCoord Coordinate {
      point [
        -1.000000 -1.000000 -1.000000,
        0.000000 -1.000000 -1.000000,
        ...
        0.000000 1.000000 1.000000,
        1.000000 1.000000 1.000000
      ]
    }
    vertexLink [
      1 1 1, 32
      2 2
      ...
      2 2
      2 2
    ]
    vertexOffset [
      0.500000
      0.200000
      ...
      0.200000
      0.500000
    ]
  }
}
```

4.3.5 Particle systems

4.3.5.1 Particles

4.3.5.1.1 Node interface

Particles { #%NDT=SF3DNode

exposedField	SFFloat	creationRate	500
exposedField	SFFloat	creationRateVariation	0
exposedField	SFFloat	emitAlpha	1.0
exposedField	SFColor	emitColor	1 1 1
exposedField	SFColor	emitColorVariation	0 0 0
exposedField	SFVec3f	emitterPosition	0 3 0
exposedField	SFVec3f	emitVelocity	0 0 0
exposedField	SFVec3f	emitVelocityVariation	1 1 1
exposedField	SFBool	enabled	TRUE

```

exposedField SFFloat   fadeAlpha           1.0
exposedField SFCOLOR   fadeColor           0.25 0.25 0.25
exposedField SFFloat   fadeRate            0.25
exposedField SFVec3f   force                0 -9.8 0
exposedField MFInfluenceNode influences []
exposedField SFWorldNode init             NULL
exposedField SFTIME    maxLifeTime         5
exposedField SFFloat   maxLifeTimeVariation 0
exposedField SFInt32   maxParticles        500
exposedField SFFloat   minRange            1
exposedField SFFloat   maxRange            -1
exposedField SFFloat   particleRadius      0.1
exposedField SFFloat   particleRadiusVariation 0
exposedField SFFloat   particleRadiusRate  0
exposedField SFWorldNode primitive        NULL
exposedField SFInt32   primitiveType       2
}

```

4.3.5.1.2 Functionality and semantics

The **creationRate** field defines the rate, in number of primitives per second, at which new primitives are created.

The **creationRateVariation** defines the variation in **creationRate** in number of primitives per second. The **emitAlpha** field defines the transparency of the primitive at creation time.

The **emitColor** is defines the primitive color at creation time.

The **emitColorVariation** defines the variation in **emitColor**.

The **emitterPosition** defines the position at which new primitives are created, or, in the case that a particle initializer node exists, the position of this node.

The **emitVelocity** defines the velocity of the primitives at creation time in units per second in x, y and z direction respectively.

The **emitVelocityVariation** defines the variation in **emitVelocity** in units per second in x, y and z component respectively.

The **enabled** field specifies whether the particle system should emit new primitives or not.

The **fadeAlpha** field defines the transparency of the primitive at the time the primitive ceases to exist, if the **fadeRate** is sufficiently high compared to the lifetime of the primitive.

The **fadeColor** field defines the expected color of the primitive at the time the primitive ceases to exist, if the **fadeRate** is sufficiently high compared to the lifetime of the primitive.

The **fadeRate** field defines the fade factor defined as a factor per second used to compute the interpolation between the primitive's initial transparency and color values to the expected transparency and color at the time the primitive ceases to exist.

The **force** field specifies a constant external force that influences the primitives.

The **influences** field contains nodes that influence the primitives during the primitives' lifetime.

The **init** field contains a node which describes how new primitives are emitted.

The **maxLifeTime** field defines the expected lifetime of a primitive in number seconds.

The **maxLifetimeVariation** field defines the variation in the lifetime of a primitive in number of seconds.

The **maxParticles** field defines the maximum number of primitives that can exist at any time instance.

The **minRange** and **maxRange** fields define the visibility range of the generated primitives. If **maxRange** < **minRange**, the visibility range is set to the visibility range defined in the scene.

The **particleRadius** field defines the radius of the primitives created.

The **particleRadiusVariation** field defines the variation in the initial radius of the primitives.

The **particleRadiusRate** field defines the rate at which the radius of the primitives changes in units per second.

The **primitive** field contains a node which type depends on the **primitiveType** field.

The **primitiveType** field defines the type of primitives created by the **Particles** node. The following types are defined:

- 0: point - A point represents each primitive. The **primitive** field is ignored.
- 1: line - A line represents each primitive, and is defined by the primitive's old and new position. The **primitive** field is ignored.
- 2: quadrangle - Each primitive is represented by a screen-aligned quadrangle. The **primitive** field can contain a node of node type **SFTextureNode**. If such a node exists, the texture is mapped 1:1 on each quadrangle. Any other node type will be ignored.
- 3: shape - The primitive field will contain a **SF3DNode**. For each primitive the node is displayed using the current primitive position, and scaled according to the radius value. Alpha and color values are not applied (but used for death calculations).

Each internally generated primitive store the following attributes:

- **SFFloat timeOfLife** the current time of life
- **SFFloat timeOfDeath** the time of death
- **SFVec3f position** the current position
- **SFVec3f velocity** the current velocity
- **SFFloat radius** the current radius
- **SFColor initialColor** the initial color
- **SFColor color** the current color
- **SFFloat alpha** the alpha value

The initial attribute values of a particle are calculated when the particle is emitted as follows: (assuming **Rand()** return a random float between -1 and 1 and **Rand3()** returns a **SFVec3f** where each component is a random float between -1 and 1)

timeOfLife = 0.

timeOfDeath = **maxLifetime** + **Rand ()*maxLifetimeVariation**;

position = **emitterPosition** + an offset calculated by the **init** node (if any)

velocity = **emitVelocity** + **Rand3()*****emitVelocityVariation**;

radius = **particleRadius** + **Rand ()*****particleRadiusVariation** (clamped downward by 0)

color = **initialColor** = **emitColor** + **Rand3()*****emitColorVariation** (clamped between 0 0 0 and 1 1 1)

alpha = **emitAlpha**;

Note that for velocity and color equations, a component-by-component multiplication is used.

Then at each update the values are recalculated as follows (setting **tDelta** as the time interval between updates):

timeOfLife += **tDelta**;

position += **tDelta*** **velocity**;

velocity += **tDelta***(**gravity** + sum of gravitational forces from **SFInfluenceNodes**.) **SFInfluenceNodes** may also change the velocity directly, and possibly absorb the primitive.

radius += **tDelta*****particleRadiusRate**

color = a color interpolated in HSV space between the **initialColor** and the **fadeColor** using **TimeOfLife*****fadeRate** as the interpolation fraction (clamped between 0 and 1).

alpha = **emitAlpha***(1-f) + **fadeAlpha***f where f = **TimeOfLife*****fadeRate** (clamped between 0 and 1)

Note: The above calculation of the position is a simple numerical solution to a differential equation. It is correct if the step size (**tDelta**) tends toward 0. If the **tDelta** is big the numerical error also grows large. More advanced numerical methods may be used in order to reduce this error.

Primitives are destroyed when:

- **timeOfLife** exceeds **timeOfDeath**
- **alpha** reaches 0
- the primitive is absorbed by an **SFInfluenceNode**

4.3.5.2 ParticleInitBox

4.3.5.2.1 Node interface

```
ParticleInitBox { #%NDT=SFParticleInitializerNode
    exposedField SFVec3f    size    1 1 1
    exposedField SFFloat    falloff  0
}
```

4.3.5.2.2 Functionality and semantics

The **ParticleInitBox** node describes how primitives are emitted from a **Particles** node. The **size** field defines the size of the initialization box. The **falloff** field defines the degree of randomness for generated values in the box as follows: **particlePosition** = **size*****Rand3()**^(1+falloff*128)

4.3.5.3 PlanarObstacle

4.3.5.3.1 Node interface

```
PlanarObstacle { # % NDT = SFInfluenceNode
    exposedField SFVec3f    distance  0 0 0
    exposedField SFVec3f    normal    0 1 0
    exposedField SFFloat    reflection 0
    exposedField SFFloat    absorption 0
}
```

4.3.5.3.2 Functionality and semantics

distance and **normal** define the equation of the obstacle plane. The **reflection** field defines the impact the obstacle has on a particle. The **absorption** field defines the probability of a particle dying at impact with the plane. A particle's new velocity after hitting the plane if not absorbed is defined as follows:

velocity = **velocity** - **normal***(**velocity*** **normal**)*(1+ **reflection**)

4.3.5.4 PointAttractor

4.3.5.4.1 Node interface

```
PointAttractor { # % NDT = SFInfluenceNode
    exposedField SFFloat    outerRadius 100
    exposedField SFFloat    innerRadius 10
    exposedField SFVec3f    position  0 0 0
    exposedField SFFloat    rate       1
}
```

4.3.5.4.2 Functionality and semantics

If a primitive emitted from a **Particles** node is inside the sphere given by **position** and **outerRadius**, the primitive is attracted and its **acceleration** = **rate***(**position** - **particlePosition**) / |**position** - **particlePosition**|³. This **acceleration** vector is applied to the particle in the same way as the gravity vector in **Particles**. If the primitive is in the sphere given by **innerRadius**, the primitive is absorbed.

4.4 Solid representation

Solid representation includes 3 geometry nodes (Implicit, Quadric and SolidRep) and extensions of the script language for implementation of the Arithmetic of Forms.

4.4.1 Solid modeling nodes

4.4.1.1 Implicit

4.4.1.1.1 Node interface

```
Implicit { # % NDT = SFGeometryNode
    exposedField SFVec3f    bboxSize  2.0 2.0 2.0
    exposedField MFFloat    c          [ ] # 4 coeffs for hyperplane
                                         #10 coeffs for quadrics,
                                         # 35 coeffs for quartics
    exposedField SFBBool    solid      FALSE
    exposedField SFBBool    dual       FALSE
    exposedField MFloat32   densities  [ ] # ( 1, 2)
}
```

4.4.1.1.2 Functionality and semantics

The **Implicit** geometry node defines an algebraic surface represented by an implicit equation, which makes it possible to know if an unspecified point of space is inside, on or outside the volume delimited by this surface.

bboxSize : bounding box dimension

c : coefficients of following polynomials as follows:

The implicit equation that defines an hyperplane, is as follows:

$$c_0 X_0 + c_1 X_1 + c_2 X_2 + c_3 X_3 = 0$$

The implicit equation of the second degree that defines quadrics, is as follows:

$$c_0 X_0^2 + c_1 X_0 X_1 + c_2 X_1^2 + c_3 X_0 X_2 + c_4 X_1 X_2 + c_5 X_2^2 + c_6 X_0 X_3 + c_7 X_1 X_3 + c_8 X_2 X_3 + c_9 X_3^2 = 0$$

The implicit equation of the fourth degree that defines quartics is as follows:

$$c_0 X_0^4 + c_1 X_0^3 X_1 + \dots + c_4 X_1^4 + c_5 X_0^3 X_2 + \dots + c_{14} X_2^4 + c_{15} X_0^3 X_3 + \dots + c_{34} X_3^4 = 0$$

In the above equations, the point coordinates were made homogeneous by the addition of a fourth coordinate X_3 .

solid indicates if the geometry is solid (TRUE) or is a surface (FALSE)

dual indicates that a unary duality operation has to be applied to volume initialization. Implies **solid** == TRUE.

densities has 2 values: the first value is the density code associated to the skin and the second value is the density associated to the inside.

The surface is only created when the node is displayed allowing its inclusion into a **SolidRep** node. The origin of the axis system is the bounding box's center. The surface will be clipped by the bounding box. If **solid** is set to TRUE, the volume will be closed and intersected with the bounding box.

If the **dual** field is TRUE, the volume is set to **solid** and the difference operation with bounding box will be carried out (bounding box – implicit solid).

The use of bounding boxes for **Implicit** nodes makes it possible to limit infinite surfaces.

The densities are used when including the primitive into a **SolidRep** node with solid operation. The density value is an Int32 and can be encoded.

As a benefit of the implicit definition, the function *isOutside (Point &p)* will return 3 values:

- 1 : if the point is outside the volume defined by the surface
- 0 : if the point is on the surface
- 1 : if the point is inside the volume defined by the surface

4.4.1.2 Quadric

4.4.1.2.1 Node interface

```
Quadric { #%NDT=SFGeometryNode
  exposedField SFVec3f  bboxSize  2 2 2
```

```

    exposedField SFVec4f P0      -1 0 0 1
    exposedField SFVec4f P1      1 0 0 1
    exposedField SFVec4f P2      0 1 0 0
    exposedField SFVec4f P3      0 0 1 0
    exposedField SFVec4f P4      0 1 0 1
    exposedField SFVec4f P5      0 0 1 1
    exposedField SFBBool solid    FALSE
    exposedField SFBBool dual     FALSE
    exposedField MFInt32 densities [] # ( 1, 2)
}

```

4.4.1.2.2 Functionality and semantics

The **Quadric** geometry node defines a second-degree implicit surface by using 6 geometric control points.

bboxSize bounding box dimension

The 6 geometric control points of the quadric, in projective coordinates (see 4.4.1.2.2.2) are:

P0, P1 2 points tangent to the quadric

P2, P3 2 poles of the construction tetrahedron

P4, P5 2 passing points of the quadric

solid solid (TRUE) or surface (FALSE)

dual unary duality operation to be applied to volume initialization. Implies solid == TRUE.

densities 2 values: the first value is the density code associated to the skin and the second value is the density associated to the inside.

Each point is defined using homogeneous coordinates allowing the point to be sent to the infinity. The values are relative to the unitary bounding box (from -1 to $+1$). If the absolute value of a coordinate is greater than 1, the point is outside the bounding box.

This node creates an **Implicit** node of the second degree through a geometric interface. The polynomial coefficients will be calculated according to the geometrical construction method using a construction tetrahedron and 2 passing points as described in the Quadric's construction mechanism (see subclause 4.4.1.2.2.1).

The quadric surface is only created when the node is displayed allowing its inclusion into a **SolidRep** node. The origin of the axis system is the bounding box's center. The surface will be clipped by the bounding box. If **solid** is set to TRUE, the volume will be closed and intersected with the bounding box.

If the **dual** field is TRUE, the volume is set to **solid** and the difference operation with the bounding box will be carried out (bounding box – implicit solid).

A continuous volume deformation is implemented by moving the control points.

The densities are used when including the primitive into a **SolidRep** node with solid operation (see subclause 4.4.1.3). The density value is encoded as a 32-bit integer.

4.4.1.2.2.1 Quadric's Construction Mechanism

Based on the principle of projective geometry, a geometrical construction mechanism for quadrics has been defined in [68], [70], [58] according to Pascal's Theorem.

In fact, the geometric control of any quadric goes via the extension to 3D of the principle for constructing conics, using a construction triangle and a passing point. The quadric will be constructed and controlled with the help of a "construction tetrahedron" and two passing points located in the planes defined by two particular faces of the tetrahedron. A set of six points will thus allow us to define the quadric from two conic sections that intersect in space at the two contact points on the tangent planes.




Figure 24 — Quadric's 6 geometric control points

The sphere in the above figure is constructed with the help of two circular sections belonging to two perpendicular planes. The two points P2 and P3 that serve as poles have been sent to infinity respectively along the Y-axis and the Z-axis (fourth coordinate equal to 0).

In general, for any quadric - projective, affine or metric - we shall define just six control points, P0 to P5, such that P0 and P1 are two contact points on two tangent planes to the quadric, P2 and P3 are two poles of the quadric and therefore belong to the two tangent planes, and P4 and P5 are two passing points for the quadric defining two conic sections of the quadric in the planes P0-P1-P2 and P0-P1-P3.

It is possible to determine the 10 coefficients of the quadric's general equation from the previously described quadric's geometrical construction system. (Reference: [68] – pages 100 – 114) .

4.4.1.2.2.2 Projective Coordinates System

In order to define all the quadrics with the 6 control points, it is imperative to remain in projective space. The coordinates system used in projective space is that of Grassmann-Plücker (one refers then to homogeneous coordinates), in which a point, a line and a plane respectively have 4, 6 and 4 coordinates. It can be interesting to make a parallel with non-homogeneous Euclidean coordinates. In this respect, one uses a representation of the point in 3 dimensions of the form (X Y Z). Adding to this triplet a fourth coordinate equal to 1 does not modify in any way the correct operation of all geometrical procedures known as Euclidean. The first 3 coordinates will indifferently represent integers or decimal numbers, provided that the fourth coordinate remains equal to 1. By particularizing the fourth coordinate and giving it a 0 value, one goes from a Euclidean coordinates system to an affine coordinates system, in which it becomes possible to control and use the points to infinity. Thus, an observer located in (1 1 1 0) will be positioned at infinity in the direction determined by its Euclidean position (1 1 1 1) and the origin of the coordinates system (0 0 0 1). In addition to the possibility of managing infinity in a natural way, this approach makes it possible to solve the problem of representation and calculation of rational coordinates, which is fundamental to obtain perfect precision whenever the coordinates are initially rational or brought about to become rational. Let us consider for example the homogeneous coordinates point (1/3 5/6 0 1); it will be represented by the approximate quadruplet (0.333 0.833 0 1) or, much better, by the quadruplet (2/6 5/6 0/6 6/6) (2: 5: 0: 6).

4.4.1.3 SolidRep

4.4.1.3.1 Node interface

```
SolidRep { #%NDT=SFGeometryNode
    exposedField SFVec3f    bboxSize  2.0 2.0 2.0
    exposedField SF3DNode  solidTree  NULL
    exposedField MFInt32   densityList[]
}
```

4.4.1.3.2 Functionality and semantics

The **SolidRep** geometry node holds the solid tree resulting of solid operations on solid primitives and/or others **SolidRep** nodes, and/or other BIFS 3D geometrical nodes.

bboxSize contains the bounding box dimensions that will be used for clipping.

solidTree contains the geometry to be initialized as solid. The following nodes may be included under a **solidTree** field: **Group**, **Transform**, **Shape**, **Implicit**, **Quadric**, **SolidRep**, **Sphere**, **Cone**, **Cylinder**, **IndexedFaceSet**, and **Box**.

The **solidTree** field can also be modified by solid operations applied by a script.

densityList: set of densities to select for display . If the list is empty, the default value will be **all densities**.

A **SolidRep** node can be included in a **Transform** node. The implicit operation of the children of a **Transform** node or of any grouping node is the ternary union operation.

Texture mapping onto implicit surfaces can be done using any algorithm. The gradient algorithm [97] is an example of an algorithm well suited for implicit surfaces.

4.4.2 Scripting extensions for solid modeling operations

The result of a solid operation is always preserved in the **solidTree** field of a **SolidRep** node. The **SolidRep** node is always the root of a solid tree. This solid tree is not processed until the display request. Only the root of the complete tree will be processed (and not the sub-trees).

These operations are accessed through script language, whose selected operators (+, -...) were overloaded for operations between **Geometry** nodes.

A non-implicit solid geometrical node is first of all transformed into a **SolidRep** node before being used as a solid operation's operand. If this geometry does not define a closed volume, it will not be considered.

The syntax of the solid operation is as follows:

$$\text{SolidRep.solidTree} = \{\text{Geometry} / \text{Int32}\} < \text{op} > \{\text{SolidRep}/\text{Implicit}/\text{Quadric}\}$$

The operands are:

- An implicit primitive (**Implicit** or **Quadric**) with “**solid** = TRUE” and densities;
- Or a **SolidRep** containing a sub-solid tree;
- Or a non-implicit geometry (only solid primitives are supported such as **Sphere**, **Cylinder**, **Cone**, **Box** or closed shapes defined by **IndexedFaceSet**);

- Or an 32-bit integer (**Int32**) value representing the density of the entire space (used as operand for implicit filtering of densities)

An only-assignment expression makes it possible to convert a non-implicit geometrical primitive into a bound implicit solid.

The complete set of solid operations **<op>** is composed of general arithmetic operators, logic operators (ternary logic) and filtering operators. A left to right evaluation of the expression is applied. The parenthesis can be used for grouping with precedence on the left to right evaluation.

4.4.2.1 General arithmetic operators

The operators below achieve combinations of solid forms by isolating or selecting certain regions in a combinatorial way by means of their densities. The densities are always positive integer.

Table 2 — General arithmetic operators

Op	Description	Rules for each P(X1,X2,X3,X4)	Syntax
+	Arithmetic addition of the density of two forms	$d_r(P) = d_0(P) + d_1(P)$	F0 + F1
*	Arithmetic multiplication of the density of two forms	$d_r(P) = d_0(P) * d_1(P)$	F0 * F1
-	The positive difference of two forms <i>F0</i> and <i>F1</i> .	Returns the difference between the densities if this result is nonnegative, and 0 if the result is negative.	F0 - F1
pow	Exponentiation of forms	Raises one form to the power of another form. The density of a point with respect to <i>F1</i> serves as exponent to the density of the same point with respect to <i>F0</i> .	pow(F0, F1)
gcd	Greatest common divisor	Returns the <i>gcd</i> of the densities of two forms.	gcd (F0, F1)
lcm	Least common multiple	Returns the <i>lcm</i> of the densities of two forms.	lcm (F0, F1)
%	Integral remainder	The Integral remainder operator calculates the integer remainder when the density of a point with respect to <i>F0</i> is divided by the density of that point with respect to <i>F1</i>	F0 % F1
abs	Absolute difference	Returns the result of the subtraction between the densities of two forms when this result is nonnegative. Otherwise, the return is the result opposed value.	Abs (F0 - F1)
icurt	Integral cube root	Returns the integral cube root of the density of the form.	icurt F0

isqrt	Integral square root	Returns the integral square root of the density of the form.	isqrt F0
max	Maximum	This operator is equivalent to the ternary union for n-ary logic	max(F0 , F1)
min	Minimum	This operator is equivalent to the ternary intersection for n-ary logic.	min(F0 ,F1)

4.4.2.2 Logic operators (Ternary logic)

The forms are coded with ternary logic: the density 0 for the outside, the density 1 for the skin and the density 2 for the inside.

Note: if used on an n-ary logic, the densities are taken modulo 2 to convert the value into the ternary logic.

Table 3 — Ternary logic operators

Op	Description	Rules for each P(X1,X2,X3,X4)	Syntax																
	Ternary union of two forms.	F1 <table><tr><td> </td><td>0</td><td>1</td><td>2</td></tr><tr><td>0</td><td>0</td><td>1</td><td>2</td></tr><tr><td>1</td><td>1</td><td>1</td><td>2</td></tr><tr><td>2</td><td>2</td><td>2</td><td>2</td></tr></table> F0		0	1	2	0	0	1	2	1	1	1	2	2	2	2	2	F0 F1
	0	1	2																
0	0	1	2																
1	1	1	2																
2	2	2	2																
&	Ternary intersection	F1 <table><tr><td>&</td><td>0</td><td>1</td><td>2</td></tr><tr><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>1</td><td>0</td><td>1</td><td>1</td></tr><tr><td>2</td><td>0</td><td>1</td><td>2</td></tr></table> F0	&	0	1	2	0	0	0	0	1	0	1	1	2	0	1	2	F0 & F1
&	0	1	2																
0	0	0	0																
1	0	1	1																
2	0	1	2																
>>	Ternary implication	F1 <table><tr><td>>></td><td>0</td><td>1</td><td>2</td></tr><tr><td>0</td><td>2</td><td>2</td><td>2</td></tr><tr><td>1</td><td>1</td><td>1</td><td>2</td></tr><tr><td>2</td><td>0</td><td>1</td><td>2</td></tr></table> F0	>>	0	1	2	0	2	2	2	1	1	1	2	2	0	1	2	F0 >> F1
>>	0	1	2																
0	2	2	2																
1	1	1	2																
2	0	1	2																
<<	Reciprocal Ternary implication	F1 <table><tr><td><<</td><td>0</td><td>1</td><td>2</td></tr></table>	<<	0	1	2	F0 << F1												
<<	0	1	2																

		F0	<table><tr><td>0</td><td>2</td><td>1</td><td>0</td></tr><tr><td>1</td><td>1</td><td>1</td><td>1</td></tr><tr><td>2</td><td>0</td><td>1</td><td>2</td></tr></table>	0	2	1	0	1	1	1	1	2	0	1	2	
0	2	1	0													
1	1	1	1													
2	0	1	2													
!	Ternary dual of the volume	F0	<table><tr><td>!</td><td></td></tr><tr><td>0</td><td>2</td></tr><tr><td>1</td><td>1</td></tr><tr><td>2</td><td>0</td></tr></table>	!		0	2	1	1	2	0	! (F0)				
!																
0	2															
1	1															
2	0															

4.4.2.3 Filtering/test operators

In addition a set of test functions can be applied on **SolidRep** nodes. These functions are used to filter densities while keeping the filtering inside the solid tree instead of defining the densities to be considered during display time.

Table 4 — Filtering and test operators

Filter	Description	Rules	Syntax
==	Equality filter	The density of the volume that checks the equality test, and 0 otherwise.	F0 == F1
>=	greater than or equal filter	The density of the second volume if the density of the first one is greater than or equal to the density of the second volume, and 0 otherwise.	F0 >= F1
>	Greater than filter	The density of the second volume if the density of the first one is greater than the density of the second volume, and 0 otherwise.	F0 > F1
<=	Less than or equal filter	The density of the second volume if the density of the first one is less than or equal to the density of the second volume, and 0 otherwise.	F0 <= F1
<	Less than filter	The density of the second volume if the density of the first one is less than the density of the second volume, and 0 otherwise.	F0 < F1
even	Even filter	The volume density if the density is even, and 0 otherwise	even F0
odd	Odd filter	The result is the volume density if the density is odd, and 0 otherwise.	odd F0
!=	Difference filter	The result is the second volume density if the densities are different, and 0 otherwise.	F0 != F1

Scripts are used through **Script** node. A **Script** node can be included as descending from any grouping node but is independent of the current coordinate system.