

mental ray® Functional Overview

White Paper

Document version 1.4
July 27, 2006

Copyright Information

Copyright © 1986-2006 mental images GmbH, Berlin, Germany.

All rights reserved.

This document is protected under copyright law. The contents of this document may not be translated, copied or duplicated in any form, in whole or in part, without the express written permission of mental images GmbH.

The information contained in this document is subject to change without notice. mental images GmbH and its employees shall not be responsible for incidental or consequential damages resulting from the use of this material or liable for technical or editorial omissions made herein.

mental images®, mental ray®, mental matter®, mental mill™, mental queue™, mental q™, mental world™, mental map™, mental earth™, mental mesh™, mental™, Reality™, RealityServer®, RealityPlayer®, RealityDesigner®, MetaSL™, Meta™, Meta Shading™, Meta Node™, Phenomenon™, Phenomena™, Phenomenon Creator™, Phenomenon Editor™, Phenomill™, Phenograph™, neuray™, iray®, imatter®, Cybernator™, 3D Cybernator™, Shape-By-Shading™, SPM®, NRM™, and rendering imagination visible™ are trademarks or, in some countries, registered trademarks of mental images GmbH, Berlin, Germany.

All other product names mentioned in this document may be trademarks or registered trademarks of their respective companies and are hereby acknowledged.

Table of Contents

Introduction	1
Applications	1
Architecture	2
Functionality	4
Rendering	4
Color and Shading	6
Geometry	8
Integration	11
Levels of Integration	11
Interactivity	11
References	12
A. Feature List	13

Introduction

mental ray® generates images of outstanding quality and unsurpassed realism and achieves scalable performance through the exploitation of parallelism both on multiprocessor machines and across networks of machines. Its generality allows for a wide range of applications, including the creation of state of the art visual effects for movies, full length feature animations, visualization of automotive and other virtual models with accurate lighting, physically correct simulation of architectural design, as well as content creation for games.

mental ray is designed for the efficient rendering of images on multi-processor machines and networks of machines, as well as for a tight but flexible integration as a software library in 3D content creation and CAD systems.

Further information about the usage, functionality, and integration of mental ray can be found in [Driemeyer 05 a, Driemeyer 05 b].

Applications

mental ray is used in a wide range of applications. This section lists a number of typical applications:

CAD Visualization

mental ray provides photorealistic design visualization of product data in CAD and styling systems.

Comprehensive support for trimmed NURBS surfaces, surface connectivity, and hierarchical subdivision surfaces offers the geometry handling capabilities required by CAD systems. Various controllable approximation settings provide high geometric accuracy. Large geometric quantities can be handled efficiently, including support for heavy instancing of objects.

The ray tracing capabilities allow realistic views of the model in arbitrary environments, for example using car paint shaders and HDR images to visualize the realistic appearance of a car body in a given environment.

Visual Effects

mental ray is used for the creation of the most demanding visual effects in the movie industry. With its visual quality, scalability, and complete ability to customize all aspects of shading and visual elements, it offers all functionalities required for the creation of any imaginable visual effect.

Hundreds of thousands of realistic virtual characters can be included in a single final frame using mental ray's procedural, on demand, object creation features. 3D motion blur can be computed efficiently using the rasterizer rendering mode for fast, high quality first-hit rendering. This can be combined with global illumination and ray tracing simulations for achieving true virtual cinematography.

High dynamic range imaging (HDRI) and arbitrarily many frame buffers are supported throughout in order to be able to seamlessly combine rendered elements with live action shots.

Feature Animation

mental ray is equally well suited for the creation of full length feature animations. Any particular look can be achieved by using custom shader plugins.

Special geometric primitives and rendering modes for hair and fur, as well as subsurface scattering can be used to create and render visually compelling characters efficiently.

Game Creation

mental ray is a powerful tool to create realistic lighting and shading setups for games. Using the light mapping features of mental ray, very complex and physically correct shading and lighting calculations can be performed in mental ray, and the resulting information is output to texture maps which are then used on the game engine ("texture baking").

Architectural Design

The global illumination features in mental ray provide physically correct lighting simulations, and enable the measurement of the actual lighting intensities in a building in a given lighting environment. Visualizations of architectural designs are effectively indistinguishable from photographs due to the correct simulation of direct and indirect illumination.

Lighting Design

Light profiles, spectral rendering with a larger number of color frequency samples, and color spaces which are visually more precise than RGB, can be used to correctly simulate the optical properties of virtual models.

Visualization

mental ray can equally well be exploited in other visualization applications. Support for volumetric effects and shading can be applied e.g. in fluid flow, seismic data, or medical visualization.

Architecture

Dataflow Architecture

The architecture of mental ray is founded on a network transparent scene and rendering data base. Data are loaded and cached only as needed, creating data on demand whenever possible. This allows mental ray to render scenes with great complexity even on machines with a small amount of memory.

Based on the database, a job system manages creation of all data and dependencies among jobs and database elements. Jobs include rendered image tiles, geometry tessellations, but also arbitrary other auxiliary data.

Incremental changes can be used to render multiple frames interactively, where only the changes in the scene between frames are passed on to mental ray, reusing other components without the need for recomputations.

Parallelism

The efficient support for multithreaded and network rendering in mental ray is based on sophisticated, advanced parallel rendering algorithms.

Multithreading takes full advantage of multiple processors, hyperthreaded and multi-core processors available in a system, without requiring any additional licenses. Due to the thread-safe database and the

job system distributing jobs to multiple threads, the performance of mental ray scales with the number of processors available.

Network parallelism allows the use of additional machines as rendering slaves. Rendering jobs are distributed to the machines, and scene and rendering database elements are transferred on demand, achieving scalable performance over networks of machines as well.

Satellites are available as an alternative mode of network parallel rendering in many OEM applications that contain mental ray. A number of rendering slaves (satellites) are made available to the machine running the application without requiring additional licenses, which effectively increases the rendering performance of the interactive application as well as the throughput in the batch rendering mode of the application.

Functionality

mental ray offers an optimal powerful implementation of all the features traditionally expected of photorealistic rendering software, together with unique functionality not found in any other rendering software. The following sections describe rendering features, color handling and shading, and geometry processing capabilities in mental ray.

Documentation on the usage of mental ray is contained in [Driemeyer 05 a]. mental ray's functionality is described in detail in [Driemeyer 05 b].

Rendering

Ray Tracing

The software is based on a ray tracing architecture, which allows for the flexible implementation of any imaginable phenomena and lighting effects, including reflections, refractions, global illumination, and subsurface scattering.

It uses an advanced BSP tree (more precisely a kd-tree) algorithm to speed up the ray intersection calculations. This structure is built on demand and cached. It can handle very large data sets and supports scalable multithreading.

Rasterizer

Besides ray tracing, mental ray also offers other rendering methods for situations where these can produce results more efficiently.

A rasterizer is available for efficient first-hit rendering of directly visible objects and transparency. By separating visibility sampling and shading, high quality antialiasing can be provided while performing fewer expensive shading calculations (e.g. once per pixel). Motion blur can be computed with a relatively small performance impact. This method is well suited e.g. for high quality cinematographic rendering.

A second scanline first-hit rendering method is also available, which is sometimes faster on smaller scenes with relatively simple shading.

Global Illumination

Global illumination is the simulation of all light inter-reflection effects in a scene. This includes indirect illumination caused by the scattering of light, and effects such as caustics and color bleeding: if a red table is next to a white wall, the white wall gets a slightly pink tint. But if the pink tint is lacking in an image, the image looks fake, even though it might be hard to point out precisely why. Global illumination effects are subtle but are the essence of true photorealism.

Simulation of global illumination has at least two distinct uses:

- Physically accurate simulation of the illumination in an environment, for example the light distribution inside an office building.
- Visually pleasing lighting effects for applications in the entertainment industry. Here physical accuracy is not the most important aspect; the images just have to look more believable.

mental ray offers two methods to compute global illumination, which can freely be combined. An improved photon mapping technique emits small energy packets (photons) from the light sources, traces their trajectories through the scene (with reflections, refractions, and interactions with participating media), and stores the energy information in a special data structure called photon map. In the second phase, during regular rendering, global illumination is added by collecting the color intensities of nearby photons.

Photon mapping supports caustics, participating media (e.g. light beams reflected through a mirror), color bleeding, arbitrary many light bounces, and specular and glossy materials.

Photon maps can be pre-computed and saved on disk for later reuse in multiple renderings. Photons stored close to each other can optionally be merged to reduce memory consumption in scenes where a high number of photons must be shot to provide sufficient illumination.

The *final gathering* technique computes global illumination by sampling the incoming illumination at a point over the hemisphere. These computations are not performed at every surface point, but are interpolated from few nearby final gather points. These final gather points are computed adaptively according to the local surface and illumination variation. Four different final gathering modes are provided which are optimized for ease of use, fly-through animations through static environments, and full user control, respectively.

The final gather data can also be saved on disk, and refined and re-used for multiple renderings. Final gathering gives particularly good results for diffuse surfaces with a single light scattering bounce only.

Hardware Rendering

Now that modern GPUs can execute complex programs to compute the color of each shaded pixel, mental ray can utilize this power for accelerating its high quality rendering. Many elements can be rendered entirely in hardware, and are automatically combined with elements which still need to be rendered in software.

mental ray currently exploits the advanced Cg shading language by NVIDIA [Cg] and uses OpenGL to exchange data with the graphics card. Cg versions of a large number of shaders including OEM shader libraries are available. mental mill™ [mental mill] supports the mapping of MetaSL™ shaders [MetaSL] to both C++ and Cg, which removes the need to provide and maintain separate hardware versions of shaders.

The unique implementation of hardware rendering in mental ray is the only one to provide full support for all of the following state of the art features:

- Phenomena™
- high quality antialiasing
- order independent transparency
- motion blur
- soft shadows
- tiling for large image resolutions and high quality antialiasing

A description of the hardware support in mental ray is also available in [LeFrancois 05]. For more information on mental ray Phenomena please refer to the section on Phenomena below.

Light Mapping

Light mapping, also sometimes called *texture baking*, is a method for sampling an object before rendering, and storing the results for later use. The most common application samples illumination for each point of a texture image wrapped around the object, and stores that illumination in the texture. Illumination is effectively frozen into the texture, which can later be mapped onto an object in the conventional way.

The advantage is that rendering can obtain illumination quickly from the frozen illumination, instead of computing it at rendering time. This is especially valuable for indirect illumination, which takes more time to compute than direct illumination. Such textures, called light maps, are commonly used in games.

Light mapping is very flexible in mental ray, implemented in light map shaders, and writing to special output textures. They can then be used for game contents as well as the mental ray hardware rendering mode.

Multipass Rendering

The unique multipass rendering feature allows users to output and composite shading data in sub-pixel resolution with an arbitrary number of samples per pixel. The antialiasing and resolution limitations of regular compositing can be overcome this way. The compositing and merging of passes can be controlled by shaders.

Color and Shading

Shader Interface

Most aspects of shading are performed in custom shader plugins. Operations provided by shaders include:

- surface shading
- procedural textures
- customized shadows
- volumetric effects
- simulation of camera lenses
- accumulation of shading results and other data into frame buffers
- light sources
- displacement mapping
- creation of geometry and other scene elements

The interface available to shaders is extremely flexible and allows shaders to access the complete rendering database as well as application data.

Phenomena

A Phenomenon™ is an encapsulated set of cooperating shaders and shader graphs and a set of requirements. On the outside, a Phenomenon looks just like a regular shader, with a name and external parameters called interface parameters, but on the inside it packages any number of internal shaders and rendering options that all work together to create an effect that would be too complex to achieve with a single shader.

These encapsulated shader graphs facilitate the integration of custom shader networks into applications. Complex shader graphs can be presented to the user as a single entity with a small, well defined interface and set of parameters. See also [mental mill, MetaSL].

Shadows

mental ray provides highly optimized and accurate ray traced shadows. These enable the accurate sampling of soft shadows by area light sources, colored shadows, and volumetric attenuation of shadows.

Shadow maps are also available for fast generation of shadows for cases where shadow mapping techniques outperform ray traced shadows. Soft shadowing algorithms are available with shadow maps, and self intersections are avoided.

The proprietary *detail shadow map* method collects more information about shadow blocking elements of the scene, by storing multiple samples per pixel and supporting colored, transparent shadows. This technique combines advantages of ray traced shadows and shadow maps and is especially useful for shadowing of fur and hair, where regular shadow maps do not provide the necessary resolution in terms of z depth and shadow map pixels.

Textures

While arbitrary texture mapping can be implemented in shaders, mental ray's built in texturing code supports texture tiling and caching for very large amounts of texture maps, arbitrary texture formats including high dynamic range images, high quality texture filtering, bump mapping, and high frequency displacement mapping.

Volumes

Volumes and atmospheres (i.e. a volume surrounding all objects) are implemented in volume shaders. Volumes can interact with all aspects of shading, including shadows, global illumination, and volume caustics, caused by the scattering of light in a participating medium after being specularly reflected or refracted. Volumes can be nested.

Frame Buffers and HDRI

Colors are internally stored in frame buffers which can be represented in a number of available formats, including various high dynamic range (HDR) formats. Arbitrarily many frame buffers can be generated in order to output multiple passes for compositing, or to output additional information.

Frame buffers are stored on disk during rendering in order to use arbitrarily many and arbitrarily large frame buffers without running out of memory. This is implemented using memory mapping.

Textures and output images can be represented in more than 30 common image formats, including HDR formats like OpenEXR.

MetaSL and mental mill

Shaders in the shading language *MetaSL*[™], and the *mental mill*[™] technology are supported by mental ray [MetaSL, mental mill]. Shaders provided in MetaSL can be rendered seamlessly in software and in hardware, and can be adapted to the particular purpose independently of their origin.

mental mill provides a way to create Phenomena graphically. This solution makes it practical to use complex shader graphs containing very simple nodes called Metanodes, since the networks are combined into larger MetaSL shaders at creation time and thus avoids the run time overheads of shader connections and large number of shader nodes.

Light Profiles

mental ray supports IES and Eulumdat light profiles. Light profiles are made available by vendors of physical lamps to describe precisely how much light is emitted in which direction. Light profiles are generated by measuring emitted light at the points of a grid placed around the light. For any light direction, this determines the light energy by interpolating the nearest grid points. Eulumdat is a European format while IES is an international standard.

Color Spaces and Spectral Rendering

RGB colors do not provide sufficient color frequency resolution for all applications, like lighting simulations or wavelength dependent refractions. Color calculations can therefore be performed in mental ray in a number of other, more accurate tristimulus color spaces, like CIE XYZ.

If this is not sufficient, color can also be represented by spectra, i.e. a number of separate frequency samples, allowing wavelength dependent optics with high accuracy.

BRDFs

Actual measured surface characteristics can be represented in bidirectional reflectance distribution functions (BRDFs). Libraries of surfaces can be provided in this representation, and BRDFs can also be used to manipulate surface characteristics, create artificial materials, or even to represent precomputed global illumination characteristics. mental ray provides data formats, conversions, and operations on BRDFs which can be used by BRDF shaders.

Contour Rendering

Contour lines can be an important visual cue to distinguish objects and accentuate their forms and spatial relationship. Contour lines are especially useful for cartoon style animation productions.

Contours can be placed at discontinuities of depth or surface orientation, between different materials, or where the color contrast is high. The contour lines are anti-aliased, and there can be several levels of contours created by reflection or seen through semitransparent materials.

Geometry

Free Form Surfaces

mental ray supports free-form curves and surfaces in non-uniform rational B-spline (NURBS), Bézier,

Taylor (monomial), cardinal or general basis matrix form. Any of these forms may be rational and may be of degree up to 21. Surfaces may be trimmed.

Polygons

Support for polygonal geometry includes efficient representations of triangle lists, arbitrary concave polygons, and polygons with holes.

Hierarchical Subdivision Surfaces

Hierarchical subdivision surfaces are available in mental ray in conjunction with the the *mental matter*[®] technology [mental matter]. It implements the Loop subdivision scheme for triangle meshes, and the Catmull-Clark subdivision scheme for quadrilateral meshes. Meshes of polygons with arbitrarily many polygons can also be rendered using a Catmull-Clark subdivision.

Further features include a C++ API for the construction and manipulation of surfaces in geometry shaders, crease edges with fractional sharpness, trimming, adaptive local refinement, multiresolution editing with wavelet decomposition and variable data-dependent smoothing, and conversion of subdivision surfaces to NURBS and of NURBS to subdivision surfaces.

Approximations

Surfaces and polygons are internally tessellated to triangles. This tessellation can be adapted to the respective requirements using a number of approximation techniques, including parametric tessellations, control by edge length, distance, and curvature (angle) criteria, view dependent tessellations to pixel level or below, and so called fine approximations which work memory-efficiently with large surfaces by tessellating them in smaller portions. Fine approximations work particularly well with detailed tessellations of high frequency displacement maps.

Instancing

Geometric objects, cameras, lights, and instance groups can be instanced multiple times in a scene. Heavy instancing provides for rendering of many copies of an object at different transformations without replicating the object data in memory.

Displacement Mapping

Surfaces and polygons can be displaced along the normal using displacement maps controlled by shaders. mental ray takes care of tessellating the surfaces finely enough and with memory efficiency in the presence of high frequency displacement maps.

Hair and Fur

Special hair primitives are provided to render hair and fur efficiently. The primitives enable the creation and rendering of large amounts of hair efficiently from spline curves. Specialized ray tracing acceleration data structures are employed in this case. Large amounts of hair can also be rendered efficiently using the rasterizer and detail shadowmaps, keeping render times and memory consumption low while providing good antialiasing of the hairs as well as optional efficient motion blur.

Geometry Shaders and Placeholders

Geometry shaders are C or C++ plugins that procedurally create geometric objects and other scene elements. Objects can be created on demand with geometry shaders when a given bounding box is hit.

In addition, placeholder objects allow for the loading of objects on demand from disk just in time when a specified bounding box is hit.

Also, large objects can optionally be deleted from memory during rendering and reloaded on demand from an .mi input file when needed.

Integration

Levels of Integration

mental ray can be integrated into digital content creation and design and engineering tools in three different ways, offering different levels of tightness and interactivity:

1. mental ray is available in a C++ library version for direct integration into the application. A comprehensive API allows for the control of all rendering aspects, and to pass all elements directly to mental ray. This is the tightest integration possible, and allows full scene database sharing across the application, the mental ray library, and shader plugins. The integration API is described in detail in [Driemeyer 05 b].
2. A translator application can be provided which reads the database of the content creation system and passes it to the mental ray library.
This is a library integration like 1.), but does not necessarily provide optimal interactive rendering in an application. It allows for example the batch mode rendering of 3D content which is present in the proprietary format of the content creation tool.
3. A translator plugin exports the scene in mental images' .mi scene file format, which is then rendered using the standalone version of mental ray.
This is very useful for example for the operation of a render farm. It also enables the transfer or further processing of the scene data for rendering. This method is sometimes more efficient than 1.) and 2.) since no application data needs to be loaded into memory, and hence more memory may be available for rendering.
mi export mode is a byproduct of the integrations in 1.) and 2.) and can also be implemented just by writing an exporter from the application.

Interactivity

Integrations of type 1.) in particular result in rendering systems characteristic of a very high interactivity. Advantages of a tight integration include:

- The scene can be rendered directly in memory, without going through input and output to the file system, which greatly speeds up scene preparation.
- Interactive rendering using incremental changes.
- Scene elements can be changed interactively between frames or even during rendering of a frame.
- Shaders have full access to the application data, and the application has direct access to rendering data. Shaders can for example query the application data to affect rendering. Images and other rendering output can be fed back to the application directly in memory.
- The application and mental ray can cooperate in the efficient use of resources. If for example memory is tight in the application, it can ask mental ray to release as much memory as possible, in order to keep the application operable. If on the other hand memory is tight for rendering, mental ray can prompt the application to release memory as needed or to control how mental ray should proceed if no sufficient memory is available.
- Diagnostics can be customized to the application, filtering and modifying what is presented to the user. For example, diagnostics output can be internationalized.

- With the availability of custom shaders, mental ray can mimic precisely the shading models and parameters of the application, and can even extend it. A scene set up in the application can thus be rendered without additional user intervention, but with the added rendering capabilities of mental ray.
- Rendering within an integrated mental ray library enables rendering with the licensing and protection schemes of the application, and avoids any data security issues implied by export of the scene to the file system.
- To speed up the (interactive) rendering of the application, any number of mental ray co-processes on a network of machines can be utilized from within the application. In this setup, the integrated library within the application can be relieved of performing rendering computations, and can help reducing memory consumption on the master host.

References

- [Cg] Randima Fernando, Mark J. Kilgard, *The Cg Tutorial: The Definitive Guide to Programmable Real-Time Graphics*. Addison-Wesley, 2003.
- [Driemeyer 05 a] Thomas Driemeyer, *Rendering with mental ray*. 3rd Edition. Springer, 2005 (to appear).
- [Driemeyer 05 b] Thomas Driemeyer, Rolf Herken, *Programming mental ray*. 2nd Edition. Springer, 2005 (to appear).
- [LeFrancois 05] Martin-Karl LeFrancois, “mental ray Phenomena on the GPU”, *GPU Gems II*. Chapter 13, Addison-Wesley 2005.
- [mental matter] *mental matter Modeling Library, version 3.4: Functional Overview*. Document Version 3.5. mental images GmbH, Berlin, Germany, 2004.
- [mental mill] *Functional overview of mental mill™*. mental images GmbH, Berlin, Germany, 2005.
- [MetaSL] *Design specification for MetaSL™*. mental images GmbH, Berlin, Germany, 2005.

A. Feature List

Rendering

- ray tracing architecture for automatic, accurate, and general simulation of reflections, refractions, shadows, and complex illumination
- physically correct simulation of global illumination
 - photon mapping, simulating all possible light paths
 - computation of caustics using dedicated photon maps
 - optional merging of photons to reduce memory consumption
 - four final gathering modes for efficient and easy to use computation of one or more diffuse indirect light bounces, optionally combined with photon mapping
 - participating media: volume caustics and multiple volume scattering
- anti-aliasing
 - strictly deterministic low-discrepancy sampling
 - automatic anti-aliasing of all features such as transparency, refractions, reflections, shadows, area lights, textures, caustics and volume caustics, and global illumination and multiple volume scattering
 - adaptive recursive oversampling and under-sampling (fewer than one sample per pixel taken initially, automatically refined where necessary)
- hardware rendering, supporting:
 - Phenomena
 - Cg shaders
 - order independent transparency
 - motion blur
 - antialiasing
 - tiling for large image resolutions and high quality antialiasing
 - combination of software and hardware elements
- motion blur, depth of field
 - object (transformation) motion blur
 - per-vertex motion deformations
 - multi-segment vertex motion blur for curved motion blur
 - full motion blur of reflections, refractions, light sources, shadows, caustics, and global illumination
 - depth of field implemented in lens shader
 - motion and depth output for post processing motion blur and depth of field
- light mapping
- multipass rendering in sub pixel sample resolution

Performance

- full dataflow architecture with caching and on demand data creation
- multithreading parallelism
- network parallel rendering on heterogeneous networks of machines
- incremental changes
- rasterizer for fast, high quality first generation rendering
- scanline rendering method
- BSP tree (kd-tree) acceleration method for fast ray tracing, with memory caching for very complex scenes
- hierarchical grid acceleration method for fast ray tracing
- full use of 64-bit architectures

Color and Shading

- programmable C and C++ shaders, with full access to scene data
- shade trees and shader graphs
- encapsulated shader graphs (Phenomena)
- MetaSL shaders, both in software and hardware rendering modes
- shader types: material, texture, light, volume, shadow, atmosphere, lens, environment, photon, photon emitter, displacement, output, contour, geometry, inheritance, and state shaders
- volume shaders implement non-geometric volume effects
- built-in shaders
- arbitrary many and arbitrary large frame buffers, stored on disk during rendering
- arbitrary frame buffer and image data types, including color, depth, motion, normals, labels, and coverage
- HDRI (high dynamic range imaging) using RGBE, float, and half float data types
- over 30 image types for textures and output, including OpenEXR
- IES and Eulumdat light profiles
- color spaces other than RGB
- BRDFs
- spectral rendering
- ray marching for volume shading
- incremental image preview during rendering

- contour rendering
- subsurface scattering
- shaders:
 - extensive base shader library
 - subsurface scattering
 - glossy reflections & refractions
 - metallic car paint
 - physically correct shaders
 - spectral rendering shaders
 - contour lines
 - depth of field
 - 2D motion blur

Shadows

- ray traced shadows
- shadow maps
- detail shadow maps
- area lights: spheres, disks, rectangles, cylinders, arbitrarily shaped geometric light objects
- shadow shaders
- volume cast shadows and volume self-shadowing

Textures

- texture tiling and caching
- over 30 texture image formats
- mipmapping and pyramid textures
- environment maps
- memory-mapped textures and memory-mapped pyramid textures
- high quality texture filtering

Geometry

- free-form surface bases: Bézier, B-spline, NURBS, Cardinal, and arbitrary basis matrix, Taylor; rational and non-rational, up to degree 21
- any number of trimming curves, hole curves, and special curves and points
- connectivity between surfaces
- polygons
- hierarchical subdivision surfaces:
 - quads (Catmull-Clark scheme)
 - triangles (Loop scheme)
 - Catmull-Clark subdivision of meshes of polygons with arbitrarily many vertices
 - sharp features: creases and points
 - wavelet-based multiresolution editing and data compression/LOD (mental matter)
 - conversion from and to NURBS (mental matter)
- trimming
- approximation, controllable by parametric subdivisions, edge length, analytic distance, angle
- displacement mapping
- texture surfaces, multiple texture coordinates
- arbitrary per-primitive data in polygonal objects
- special hair primitives
- instancing
- object and instance flags to control visibility, shadows, reflections, etc. per object
- selection of lights for illumination and shadowing per instance
- on-demand loading of object files and on-demand procedural object creation using geometry shaders
- optional on-demand re-loading of large data from the .mi input file
- placeholders

Integration

- optional library form for integration into OEM products
- translators and plugins are available for a variety of front-end applications
- C language API for entire scene description language and other features
- simple, efficient, and full-featured hierarchical scene description language
- material and parameter inheritance
- OEM specific software protection and licensing technology