HYDRA USER'S GUIDE

Abstract

This document is intended to be a User's Guide to Hydra. It is aimed towards standards development engineers and other technical users who may be interested in using Hydra. Hydra is a three-dimensional extensible markup language (XML) instance viewer/editor that was developed to aid in standards development efforts. Its purpose is to allow rapid, large-scale visualization of the overall structure of an XML document, human-comprehensible searching, and editing. It uses OpenGL (http://www.opengl.org) to display the XML document as a tree structure that can be manipulated in various ways by the user.



Figure 1: The main window with an open document and one node selected

Overview

The program works by laying out an XML document as a three-dimensional tree

(mapped onto the user's two-dimensional computer screen) in which objects representing pieces of the document are arranged according to their position within the document. Each *node* (for example, each cone or cube) represents an XML element or attribute in the opened document's structure. Additional information is displayed in the tree using shapes, colors, and varying sizes and positions. Any node with a picture on it contains some data value that can be edited; any solid-color node exists only for structural purposes. Nodes can be rotated (with respect to the user, they are rotated about the x-axis). When a node is rotated, everything in that node's subtree rotates as well. The entire tree can be dragged (left, right, up, and down) and zoomed.

Conceptually, the program has two major modes: either the program is in edit mode, or it's in search mode. When a document is first opened, it appears in edit mode. Edit mode allows the user to expand and collapse parts of the tree structure, filter out some parts of the tree (in which case they are hidden, but are still included when the document is saved), and change the data inside nodes.

By toggling the "Search Mode" indicator (further description below), the user can turn on (and off again) search mode. Search mode allows the user to enter a text string search query (every time a character is added or deleted, the results are updated) and have only relevant parts of the tree displayed. When in search mode, the user works on a separate version of the document tree – that is, all the data is the same, but the view onto that data can be different between the search and edit modes.

Program Layout

The main window has four major areas (fig.1), arranged top to bottom as follows: the menu bar, the search bar, the tree display, and the information/editing area. The tree display will be discussed first, as it is necessary to understand how the tree is laid out before looking at the rest of the program.

The Tree Display

The tree display area is the central area of the program. It displays the visual representation of an XML document. The tree display works as follows:

- The user can rotate the tree, or parts of the tree, around the axis that would run along the tree's trunk.
- The "top" of the tree is at the far left. When a node is expanded, its children are shown to its right. When initially opening a document, only the root node is visible. Conceptually, the complete tree structure generally looks like a pine tree turned on its side, with the root of the XML document at the "top" (left) of the tree.
- Each node with at least one element below it in the document hierarchy (one "child") has a *connector* between itself and its children. If there is only one child, the connector consists of a cylindrical bar linking the parent and child (fig. 2). Otherwise, the connector consists of a cylindrical bar and a *child disc* (fig. 3). The child disc is a regular polygon of *n* sides, where *n* is the number of children. The children are then connected to the disc at the vertices. (For example, if there are 4 children, the parent is connected by a cylindrical bar to a square-shaped disc, with a child at each corner.) When a node is in its collapsed state, the connector shrinks,

and children are hidden (fig. 4).



Figure 2: a node with one child



Figure 3: a node with many children



Figure 4: a node with many hidden children

Mouse controls

- The left mouse button is used to select and rotate nodes. When selected, a node becomes surrounded by a translucent yellow sphere. Only one node is selected at a time.
 - Double-clicking on a node will expand or collapse it.
- Holding the middle mouse button (or pressing and holding the wheel, on many mice) while moving the mouse vertically, zooms in or out
- Holding the right mouse button while moving the mouse drags the tree left, right, up, or down, accordingly.

Keyboard shortcuts

- The up and down arrow keys move the selection to the previous and next sibling node, respectively.
- The left arrow moves the selection to the parent of the current node.
- The right arrow moves the selection to the first child of the selected node.
- F1 toggles search mode on/off.
- F2 isolates all search results.
- F3 isolates the current node.

The Menu Bar

The menu bar, as well as standard program options (Open, Save, etc.), also contains the following actions and options:

- The "File" menu
 - "Open" opens a new XML file, "Save" saves the current file, "Save As" saves the current file with a new name, and "Exit" quits the program.
 - "Save Using Stylesheet" allows the user to select an output file and an XSL stylesheet. Hydra then applies the stylesheet to the XML tree before writing to the output file. In this way, it is possible to generate HTML output, XML output in a different format, etc.
 - "Load a Filter" allows the user to select a new set of filters to apply to the currently open document. Hydra's filter files are simply a list of XPath expressions, one per line, each of which describes a node to be hidden.
 - "Clear all Filters" makes all filtered nodes visible again.
 - "Export Filter" creates a new filter file containing any filters already loaded from an input filter file, as well as nodes that have manually been filtered out during the current session.
- The "View" menu
 - "Automatic Camera Adjustment" slides the tree left, right, up, and/or down to keep the area that the user is currently working on roughly centered on the screen. This option can be toggled on or off.
 - "Show Fading Node Info" briefly displays, in the bottom-right corner of the tree display, the names of the selected node's ancestors in the document

hierarchy.

- "Increase/Decrease Node Size" incrementally increases or decreases the size of the nodes (the cones or boxes representing XML elements).
- "Increase/Decrease Node Spacing" incrementally increases or decreases the horizontal distance between parent and child nodes.
- "Reset Camera to Start Position" restores the point of view that is presented when the document was first opened.
- The "Tree" menu
 - "Isolate Selected Node" collapses everything in the tree except for the nodes that lead from the root directly to the selected node. After this action, the selected node is always at the last visible level of the tree (NOTE: when in search mode, selecting this option takes the user back to edit mode)
 - "Expand/Collapse Current Node" shows or hides the current node's children, if there are any.
 - "Show/Hide Current Node's Attributes' if the selected node's XML tag contains any XML attributes, displays them in a satellite ring just below the current node on the tree.
 - "Isolate All Search Results" performs the same action as in "Isolate Selected Node", but using all current search results as the selected nodes.
 - "Show Entire Tree" fully expands the document, showing all nodes.
 - "Show Current Node's Raw XML Data" brings up a separate window with the XML text data corresponding to the currently selected node and its subtree. Text is selectable (and therefore can be copied, for example, to another application).
 - "Filter Out the Current Node" applies a filter to the current node and its subtree making it invisible. The resulting output document, however, will still contain any invisible nodes' data.
 - "Filter Out the Current Node's Children" same as above, but applied to all children of the current node, and not the node itself.
 - "Add a New Element/Text Element" (NOTE: currently in pre-alpha state) attaches a new element or text node as a child of the currently selected node, when possible.
 - The "Options" Menu
 - "Perform Strict Validation Checking" chooses whether to perform strict validity checking against the schema. If deselected, only non-well-formed XML documents will generate error messages upon opening; if selected, documents which do not properly conform to their referenced schema(s) or DTD(s) will generate error messages upon opening as well.
 - "Use Wireframe Display" for certain machines, this may provide a slight speed increase, at the expense of readability. It does, however, look extraordinarily cool.
 - "Toggle Font Style (Bitmap/Stroke)" this tells the renderer to display fonts in one of two ways. Bitmap fonts always face the user, but also do not scale when the user zooms in or out, and are slower. Stroke fonts rotate with the tree, and zoom as expected. Stroke fonts are the default.

- "Draw Nodes as Boxes" use untextured cubes as opposed to cones to represent XML elements. (Attributes are already represented as cubes, so no change is made to the way they are drawn.) This may increase performance on some systems.
- The "Network" Menu
 - This is discussed at the end of the guide under the "Additional Features" heading.
- The "Help" Menu
 - "Show Legend" opens a new window containing a legend which shows the objects used in the tree display and their descriptions.
 - "Contents" and "About" are not yet implemented.

The Search Bar

The search bar contains, from left to right:

- The "Search Mode" toggle button, which turns search mode on or off
- The search field, which allows the user to enter a search string when search mode is on.
- Checkboxes to indicate whether to search in the XML elements' names, the XML elements' data, or both.

Information/Editing Area

This area displays the currently selected node's information.

- The upper portion of the area displays a subset of the node's ancestors (to help place the node in a context). The ancestors are printed from left to right, similar to the tree's layout, with the current node's parent being the last node on the right.
- The lower portion of the area displays the current node's tag name and data (if available). If the user has selected a text node, the data field is editable. Changes are saved automatically.

Additional Features

Filtered Nodes

- Useful for removing unnecessary clutter from the 3D view of a document.
- It is possible to filter out nodes either by loading a filter file, or by selecting one of the filter options from the Tree menu. When a node is filtered out, it disappears, along with all of its descendants. The space it filled in it's parent's connector remains, but is blank.
- When a node is filtered out, it does not appear in either the edit or the search mode trees. However, when saving the document, the filtered node's data will be included.
- Filter list files, by convention, have the suffix ".fil"

Network Connectivity

• At startup, Hydra attempts to bind to port 4096 to listen for incoming socket connections.

- If a connection is detected, Hydra can accept commands from any client which provides them as a character stream in the proper syntax
- Currently, two commands are supported
 - "isolate" will isolate the current node
 - "select|/some/x/path" will change the selection to the XML element at /some/x/path
- The "Network" menu also contains an option to connect to another running instance of Hydra. After entering a hostname or IP address and port number for the instance to which you want to connect, Hydra will send "select" and "isolate" events to the other networked instance. Currently, Hydra only supports one incoming and one outgoing connection.