TEX and LATEX on the Web via IBM **techexplorer**

Robert S. Sutor
Interactive Sci. Publishing Group
IBM T. J. Watson Research Center
P.O. Box 218, Yorktown Heights, NY 10598 USA
sutor@us.ibm.com

Samuel S. Dooley
Interactive Sci. Publishing Group
IBM T. J. Watson Research Center
P.O. Box 218, Yorktown Heights, NY 10598 USA
dooley@watson.ibm.com

Abstract

The IBM **techexplorer** Hypermedia Browser™ is an application for the interactive publication of scientific and technical documents. The original project started as an experiment at IBM Research to see how an implementation of a subset of TEX, LATEX, and AMS-LATEX could be extended to support interactive viewing of documents for a computer algebra system. This interactivity is accomplished via support for hypertext, multimedia, user-defined pop-up windows and menus, and a modular architecture that allows connections with other applications and Java applets. IBM's **techexplorer** provides an alternative to HTML-based solutions for presenting scientific and technical documents on the World Wide Web and is being used for scientific journals, and educational courseware and textbooks.

The Intord.ctory Edition of **techexplorer** operates as a Netscape Navigator Plug-In and is available for several platforms, including Windows 95/NT, Sun Solaris, and IBM AIX. In addition to being able to display full documents using the supported LATEX language, **techexplorer** also implements the new Mathematical Markup Language being prepared by the HTML Working Group of the World Wide Web Consortium. In this paper we will give an overview of **techexplorer** and detail how it can be used to deliver mathematical articles, books and course materials via the World Wide Web. Future directions regarding our plans for opening the architecture of **techexplorer** and how that relates to the authoring of scientific and technical documents for the Internet will also be discussed.

– – ∗ – –

Introduction

The World Wide Web provides one of the greatest opportunities that the publishing industry has seen in this century, as well as one of its most perplexing challenges: how to produce interactive electronic alternatives to printed textbooks and journals that take advantage of the unique characteristics of electronic media and of the Internet in a way that is both intellectually engaging and economically viable. This challenge is even more rewarding in the arena of scientific and technical publishing, where the complexities of mathematical layout and the richness of the information contained in technical documents cause special problems not encountered with other kinds of documents, but also give rise to exciting possibilities for creating truly interactive materials that are useful for distributed and distance learning, interactive courseware, and electronic journals.

However, until recently the publication of documents containing a high degree of technical content on the World Wide Web has been extremely awkward, due to the absense of HTML support for mathematical notation. While TEX and LATEX have become a widely accepted standard for publishing scientific and technical documents, authors and publishers have had no convenient way of electronically disseminating documents written in this form. As a result, we have had to make compromises in various ways as authors to allow our materials to take advantage of the possibilities of the Internet, either by using tools for converting TEX/LATEX markup into HTML, or by using static images (GIF, PDF, etc.) for mathematical notation. Such conversions result in documents of poor visual quality, that fail to adapt well to a wide range of display and printer hardware, and that fail to preserve the rich semantic information present in technical documents.

While the browser development and Internet standards communities have long acknowledged the shortcomings of HTML for the presentation of mathematical notation, early efforts to extend HTML with additional primitives to address the needs of the technical publishing community (such as HTML 3.0, as well as early versions of HTML Math) have been largely unsuccessful, due to the relatively specialized nature of mathematical notation. The approach now being used in more recent efforts is to have the Internet community support XML as a general extension mechanism for HTML, and allow the technical publishing community to define the Mathematical Markup Language (MathML) as an XML application. This alternative holds greater promise for the future, but it will be some time before this approach can be fully supported by software developers.

The IBM **techexplorer** Hypermedia Browser™ provides an alternative to HTML-based solutions by dynamically rendering a large subset of TEX and LATEX markup without converting the original document source to an alternative markup or binary format. Instead, when **techexplorer** reads a LATEX document, it parses the original document source into an internal object-oriented representation that mirrors the visual structure of the document. Such a representation enables high-quality dynamic rendering at varying screen resolutions and sizes, rapid document reflow and redisplay, and the opportunity to implement various extensions to the language that facilitate the use of hypertext and the inclusion of multimedia content.

This object-oriented approach allows **techexplorer** to support a much richer model of active mathematical content. When a mathematical expression appears in the text of a document, **techexplorer** can represent both the visual appearance of the expression and the underlying semantic mathematical content being presented by the visual notation. This capability allows **techexplorer** to support interactive mathematical manipulations within a document that operate not on the visual appearance of the document, but on the underlying structure of the mathematical expressions it contains. Using this capability, we have used **techexplorer** to develop interactive courseware that supports symbolic problem solving and intelligent graphical exploration, on a platform that allows interactive technical documents to be delivered over the Internet.

Project history

techexplorer is our second generation TEX-based hypertext system. The original application, known as HyperDoc, was developed as a viewer for documents for the computer algebra system now known as Axiom and now distributed by the Numerical Algorithms Group, Ltd. (NAG) Axiom is a sophisticated system for performing mathematical computation that offers two- and three-dimensional graphics, a hypertext help package, and various forms of output, including LATEX and Fortran. HyperDoc was developed in the late 1980s for Unix platforms, and serves as Axiom's hypertext front end for viewing text intermixed with the results of computations. The Axiom/HyperDoc link enabled users to open a workspace by clicking on Axiom input, or start a graphics manager by clicking on a graphic.

The first challenge we faced when extending HyperDoc was that it accepted a non-standard dialect of LATEX, and so our initial efforts were directed toward improving HyperDoc to support standard LATEX. However, HyperDoc still did not render many common TEX/LATEX control sequences, and what is perhaps most surprising, it had very weak support for displaying mathematics. Some early implementation choices made adding such support very difficult and hindered porting HyperDoc to other platforms such as Microsoft Windows. Although HyperDoc had excellent connectivity to Axiom and to the graphics manager, there was no way to update a document in-place with the results of a computation or with a modification to a graph. For these reasons, in 1994 we embarked on developing a completely separate implementation of a large subset of standard TEX and LATEX, with extensions to support the interactive viewing of documents.

That project's efforts resulted in an early version of **techexplorer** (then known as "Saturn") that produced familiar output from TEX/LATEX source. Our approach to orthogonally extending TEX and LATEX with control sequences for hyperlinking, multimedia, and interaction with mathematical software ensured **techexplorer**'s compatibility with TEX and LATEX markup for printed documents. An early version of this standalone edition of **techexplorer** is used as the front end for NAG's Axiom for Windows product.

Our decision to provide a Windows 3.1 and later a Windows 95/NT implementation of **techexplorer** was prompted by the sophisticated set of available tools geared toward rapid C++ code development, a rich set of user interface components, a robust implementation of interprocess communication via Object Linking and Embedding (OLE), and a potentially large user community. Our aim was to leverage these technologies and quickly develop a framework for rendering and interacting with mathematical documents.

It became evident in early 1996 that we could augment our core technology to deliver interactive scientific and technical documents over the World

Wide Web via the Netscape Navigator plug-in interface. The first version for Windows 95 was made publicly and freely available in May 1996. This "Introductory Edition" of the **techexplorer** Plug-In allows authors and publishers to effectively expand the reach of their articles, books, and journals by making them available on the Internet. As the community of **techexplorer** users began to grow, we realized that a UNIX edition of the plug-in was a high priority for our colleagues in the scientific community. In September 1997, we released our first "Preview Release" of **techexplorer** on a UNIX platform on IBM alphaWorks.

In parallel to the work on the core TeX/LaTeX technology, we realized early on that the dynamic rendering and object-oriented representation used in **techexplorer** could provide an excellent framework for an interface for active mathematical documents such as interactive textbooks. Combined with the group's earlier expertise in computer algebra system development, we felt that the creation of an electronic textbook for linear algebra would be a natural application of the **techexplorer** technology. As a result, a stand-alone version of **techexplorer** was used as the framework for an electronic version of the textbook *Linear Functions and Matrix Theory* by Bill Jacob, that will be appearing as the first volume of the forthcoming Springer Interactive CourseWare Series. This interactive textbook combines **techexplorer**'s dynamic document model with the powerful symbolic computation facilities of Axiom and with a collection of Java graphical exploration tools to allow a reader to interact with the course material in a number of novel ways.

Today, the Interactive Scientific Publishing Research Group at IBM Research continues to distribute the **techexplorer** Plug-In, Introductory Edition for Windows 95/NT, IBM AIX 4.1 and SUN Solaris 2.5, with more platforms planned for the future. The **techexplorer** product line will continue to evolve as we create new tools and technologies for the Internet delivery of scientific and technical journals, reports, textbooks, and courseware.

techexplorer overview

In creating **techexplorer**, we set out to implement a majority of the standard LaTeX control sequences and environments, as well as a substantial collection of the commonly used features from plain TeX. Support for these features, especially during the early development stages of the introductory edition of **techexplorer**, has traditionally been user and application driven. At the time of this writing, virtually all of the standard LaTeX commands are available; the **techexplorer** user guide lists the LaTeX and TeX commands that are supported, as well as the **techexplorer** extensions that have been added. In addition to the support for standard LaTeX, full support for AMS-LaTeX is planned for the near future. Parsers for subsets of SGML, XML, and MathML presentation tags have also been implemented that produce the same object-oriented representation used by the LaTeX parser. As a result, markup written in these languages can be embedded in LaTeX documents, and vice versa, and rendered using **techexplorer**.

When **techexplorer** parses a document, an internal tree structure of objects is created that represents the document contents. Thus when the user clicks on the display screen, **techexplorer** has enough information to identify the object in the structure hierarchy under the position of the mouse cursor. This means, in particular, that **techexplorer** can provide hypertext links or maintain status messages that are updated as the cursor passes over different objects in the document. Different flavors of **techexplorer** links can:

- navigate to another location, either in the current document or in a different document, possibly in a different frame;
- start an application;
- play an audio or video clip from a URL;
- pop up various kinds of dialog boxes for user input;
- display fully-formatted text in a pop-up window;
- display one or the other of two expressions; or
- send input to another application, and place any output generated as a result into the current document.

Documents delivered over the World Wide Web are easier to navigate if they are broken up into reasonably sized sections with a rich collection of hyperlinks. For many documents, this creates a natural tree hierarchy. The commands `\aboveTopic`, `\previousTopic`, and `\nextTopic` allow the reader to move up, left or right, respectively, in this hierarchy, as it is defined by the document author. When these commands are defined in a document, the default document context menu (produced by clicking the right mouse button) allows the reader to jump quickly to the corresponding sections. Commands on the document context menu also allow the reader to navigate forward and backward in the dynamic sequence of sections visited in the current session. All of these commands are also available from the **techexplorer** toolbar.

Color support is very important for high quality on-screen display and so **techexplorer** implements the commands `\color`, `\textcolor`, `\colorbox`, `\fcolorbox`, `\rgb`, `\pagecolor`, as well as the **techexplorer** extension `\colorbuttonbox`. We also provide the `\includegraphics` command to embed images in either GIF or JPEG format in a document. The `\backgroundimage` command is a **techexplorer** extension that allows the author to set the image displayed behind the text in a window.

Pop-up menus, also known as context menus, are very useful for making selections. Within an electronic version of a large textbook, for example, it is relatively straightforward to define context menus that can be used to easily navigate from section to section with the book, or within smaller sub-sections of the current section. As an example, here is a menu definition from section 3.2 of *Linear Functions and Matrix Theory*:

```
\newmenu{theorem-menu-3.2}{
    \labelLink{theorem-3}{Theorem 3}
}
\newmenu{section-menu-3.2}{
    \labelLink{sec-3.2}{Section 3.2}
    \hrule
    \usemenu{theorem-menu-3.2}{Theorems}
    \hrule
    \labelLink{sec-3.2.1}{Elementary Operations}
    \labelLink{sec-3.2.2}{The Augmented Matrix of a System}
    \labelLink{sec-3.2.3}{Equivalent Systems of Equations}
    \labelLink{sec-3.2.4}{Gaussian Elimination}
    \labelLink{sec-3.2.5}{Echelon Form Systems}
    \labelLink{sec-3.2.6}{Solutions to Echelon Form Systems}
    \labelLink{sec-3.2.7}{Row-Echelon Matrices}
    \labelLink{sec-3.2.8}{Two Remarks About Parameters}
    \hrule
    \labelLink{section-3.2.problems}{Problems for 3.2}
    \hrule
    \docLink{lfmttoc.tex}{Book Contents}
}
```

The `\hrule` commands put separator lines in the menu. Most of the menu selections are hypertext links within the section, but there is one submenu that lists the theorems in the section. (In this case, there is only one theorem.) The `\usemenu` command is used in menus to create submenus as above, but it is also used within the text to associate a menu with a particular piece of text.

An interesting and sometimes controversial subject is how **techexplorer** deals with fonts. On Windows 95 and Windows NT, **techexplorer** uses TrueType fonts directly, while under UNIX, **techexplorer** uses PostScript fonts. METAFONT fonts are not currently supported. The user can select any TrueType or PostScript font under Windows or UNIX, respectively, for use with `\rm`, `\bf`, `\it`, `\tt`, etc. Special symbols are obtained from various font collections that may be available from other sources, such as the Monotype Math fonts shipped with Lotus SmartSuite, the WordPerfect math fonts, or the Lucida fonts sold by Microsoft. For UNIX we derived several symbol fonts from the Computer Modern and AMS Symbol fonts created by BlueSky Research and Y&Y and placed in the public domain under the auspices of the American Mathematical Society. In 1998, we plan to release TrueType versions of these derived symbol fonts for use under Windows. Eventually we expect to open the font model to allow authors to map a symbol to an arbitrary character in a font of their choice.

Authoring

The Introductory Edition of the **techexplorer** Hypermedia Browser$^{\text{TM}}$ is available for free download via the IBM **techexplorer** web site.[1] This version contains all the core TEX/LATEX parsing and rendering features, and can be used to rapidly distribute scientific and technical documents on the World Wide Web. **techexplorer**'s compatibility with standard LATEX means that these documents can be developed in the most natural markup language for mathematical notation and disseminated directly without major modification and without translation into another format. The **techexplorer** extensions to LATEX provide a gentle upgrade path that allows authors to incrementally include hypertext and multimedia extensions to their documents as time and resources allow.

Several options are available for increasing the interactivity of electronic versions of LATEX documents using the **techexplorer** extensions. A **techexplorer** document can be used as an interface to other programs. Depending on the nature of the application, the integration between the document and the application can range from a model where the document is used merely to launch the external application, as in the Introductory Edition, to one where there is a much higher degree of active mathematical information being communicated from the document to the application and vice-versa, as in the upcoming Springer Interactive CourseWare.

For those authors comfortable with Java programming, Java applets can also be incorporated into LATEX documents using the **techexplorer** extensions, either by invoking stand-alone Java programs using the more general application link, or by using the **techexplorer** extension command `\javaLink` to invoke a Java method directly. In the future, we plan to expose more of the underlying **techexplorer** document model to provide a consistent interface for external Java methods to manipulate and script **techexplorer** document contents.

[1] `www.software.ibm.com/techexplorer`

Conclusions

Our primary goal for the **techexplorer** technology is to create a set of publicly available interfaces that allow **techexplorer** to serve as the "glue" for developing novel interactive scientific and technical documents. By using **techexplorer**, such documents will be able to leverage Internet programming languages, as well as specialized programs from a wide-range of scientific software applications.

Clearly, with the continued growth of the World Wide Web and increased investments in scientific and educational content, traditional scientific and technical markup languages will play a central role in online dissemination. The deployment of interactive scientific and technical documents using enhanced versions of these languages will be central to the success of the next generation of technical publishing. We in the Interactive Scientific Publishing group at IBM plan to continue enhancing the **techexplorer** products to support and to define professional quality scientific online publishing in this next generation.