

Daisy 3: A Standard for Accessible Multimedia Books

Dennis Leas
gh, LLC

Emilia Persoon
Dedicon

Neil Soiffer
Design Science

Michael Zacherle
Karlsruhe Institute of Technology

The Daisy standard for multimedia representation of books and other material is designed to facilitate technologies that foster easy navigation and synchronized multimodal presentation for people with print-reading-related disabilities.

People who are blind or have limited vision need alternative versions of books and magazines. Typically, these formats are not available locally and must be specially ordered. They are not as convenient to use as a regular book or magazine because they are large and cumbersome to store and carry around. In addition, printed material causes difficulties for people with some cognitive disabilities, such as those with dyslexia (this group is estimated to be 5 to 17 percent of the population¹) or motor impairments who can't hold a book or turn its pages. Electronic documents hold the promise of substantially increasing access to material for all of these groups.

For the blind, electronic documents are always available and can be accessed with a portable electronic device. For those with vision-related cognitive disabilities, electronic documents can be presented in a simplified

layout, with special fonts or audio synchronized with word highlighting to aid comprehension.¹⁻² Additionally, numerous hardware or software technologies, such as large buttons or eye tracking, can make documents accessible to those with motor impairments (see, for example, <http://www.webaim.org/articles/motor/assistive.php#eyetracking>).

The Digital Accessible Information System (Daisy) standard describes an open data format for the representation of interactive books that are accessible to those with print-related disabilities.³ Daisy books may have both a textual and an audio component and allow for an active reading experience. To read a Daisy book, a reader needs a hardware or software playback system. Unlike a book on a cassette tape that users typically listen to from start to finish, readers using a Daisy book player can easily move backward and forward in the book; they can move to chapters, sections, pages, or bookmarks they have created. If the reader resumes reading, the player begins playback where the reader left off. Some players allow readers to change the magnification and the colors used. Many such players also highlight what is being spoken.

One of the problems that Daisy books are particularly good at solving is image translation. For example, in technical textbooks and journals, equations, diagrams, graphs, and other image-oriented material often accompany the text. These images are not immediately accessible to blind people. For those with limited vision, images can cause problems because they often don't look clear when magnified. And for those with cognitive disabilities, cluttered images or complicated ordering within the image might hinder understanding. Some images can be made accessible by providing a textual description of the image. However, for many scientific or mathematical images, this summary information doesn't help readers understand an image's structure and relationships. Daisy books can solve this problem by using structured markup instead of an image. For mathematical expressions, Mathematical Markup Language (MathML)⁴ can represent the expressions; for diagrams and graphs, Scalable Vector Graphics (SVG)⁵ can be used. These formats let users navigate the images, hear them as spoken words, or convert them to other formats, such as braille or tactile displays.

While Daisy was designed to make materials accessible to those with print-related disabilities, the format is useful for people of all abilities. As such, the standard is an example of universal design; features such as being able to move to a section or a bookmark added by a reader are useful to everyone. The ability to have a visual, audio, or combined text-and-audio presentation generated from the same source material means that readers could choose an experience appropriate for their abilities or constraints. One example is using an audio-only version while driving a car or riding on mass transit. Because of Daisy's universal usability, several organizations have adopted the standard.

Evolution of Daisy

For the blind, braille translations of print books have provided an alternative for over a century. Braille books share many of the convenient features of print books, such as the ability to jump to pages or reread pages as desired. However, their weight, size, and binding make them hard to carry around. Braille books are large and heavy because the embossing paper is thicker than traditional paper; braille characters are approximately the same size as 36-point characters; and braille books must be loosely bound so as not to crush the braille dots. Reading a braille copy, of course, requires proficiency in braille. In some countries braille literacy has decreased significantly. For instance, US braille literacy declined significantly in the last 50 years from 50 percent of the school-aged blind children in the 1960s to 12 percent in the 1990s.⁶

In addition to the issues associated with braille, there are other factors that make a format other than traditional small-print books potentially useful to many people. For example, many older readers find it difficult to read regular print books due to normal aging processes and age-related eye diseases. Almost three million people in the US alone will suffer from macular degeneration by 2020.⁷ Audio formats, such as talking books, provide an alternative to print books. In the 1970s, analog cassette tapes dominated the talking book industry. Cassette tapes and players were widely available and relatively inexpensive, and, unlike braille, don't require any special training to use. However, audio books lack the ability for readers to easily find sections or

pages in a book. Being able to find a place in a book quickly is critical for studying or in a classroom setting where an instructor might say, "Please turn to page 107." Also, the sharing of books between organizations or persons was hindered because of incompatible formats for the recordings used in different countries.

The Swedish Library of Talking Books and Braille (TPB) started the Daisy project in 1988. One of the chief goals was to move away from analog cassette tapes and toward a digital medium. In the early 1990s, compact discs gained widespread acceptance and TPB selected this format as the target medium for Daisy distribution. A prototype Daisy player was built in 1994 and presented at the International Conference on Computers Helping People with Special Needs (ICHP) later that year. This resulted in considerable interest in the Daisy project.

In 1996, TPB invited other talking-book libraries to form a worldwide consortium to promote digital, talking books based on the Daisy ideas. In 1998, The consortium revised the format to conform to several World Wide Web Consortium (W3C) multimedia standards, such as XHTML and Synchronized Multimedia Integration Language (SMIL),⁸ as the basis of a Daisy 2.0 specification. The following year some small enhancements led to new 2.01 and 2.02 specifications.

The original goals for Daisy centered on enhancing the audio experience for listening to e-books. The first Daisy players were audio-only; the majority of players remain audio-only today. However, it was recognized early on that Daisy could serve as a format for visual players too. Those players could offer extensive functionality, such as image display, text magnification, and synchronization of text and audio. People with limited vision and learning disabilities might find such features helpful.

The Daisy 2.02 specification requires many cross-references between the textual XHTML file and the controlling SMIL file. Producing this cross-referencing by hand for real books isn't feasible. Hence, use of the Daisy format required waiting for the arrival of tools that allowed for easier production of Daisy materials. The first such professional tool was announced in 2000. Production of Daisy material began a short while after that.

Feedback from users, producers, and Daisy player manufacturers, along with a push to make Daisy a formal standard, led to a revision of 2.02 becoming the ANSI/NISO Z39.86-2002 standard. This standard was updated in 2005 to become Z39.86-2005 and is commonly referred to as Daisy 3. The standard moved from using XHTML in Daisy 2 to using XML in Daisy 3, a switch that allowed for a much richer use of elements. Daisy 3 is a functional superset of Daisy 2.

An important feature of the Daisy 3 specification is that it allows for modular extensions. The first of these extensions was for mathematics and was approved in 2007.⁹ The math extension is based on MathML, a W3C recommendation for encoding math in XML. The addition of a mathematics module to Daisy shows that Daisy can evolve by incorporating standards from other areas. Video and digital rights management are two such areas being considered. The Daisy Consortium closely monitors developments regarding XML and is continually gathering requirements for future versions.

Daisy not only is incorporating other standards, but also is being adopted by other organizations. The International Digital Publishing Forum's Open Publication Structure for e-books incorporates Daisy into its IDPF standard as its preferred vocabulary.¹⁰ IDPF includes large publishers such as McGraw-Hill, John Wiley & Sons, and HarperCollins. The National Instructional Materials Accessibility Standard (better known as NIMAS) is based on Daisy and is the legally mandated format for new K–12 textbooks and supplementary material in the US.¹¹ Extension of these requirements to higher education textbooks likely will happen in the future.

Anatomy of a Daisy 3 book

A Daisy book can consist of audio only, text and images only, or a combination of both. All book types use the same set of file types, although some of the files are optional. Unless otherwise specified in our following discussion, all file types are based on XML. The nine file types comprising a Daisy 3 book are

- package file,
- textual content file,
- image file,

- audio file,
- synchronization file,
- navigation control file,
- resource file,
- presentation style file, and
- transform file.

The package file is a required type that acts as the master file of the Daisy book during playback. It contains metadata about the book such as the language, publisher, and copyright date; a manifest listing all the files comprising the book; and a `<spine>` element that indicates the book's linear reading order.

A textual content file contains the book's text and is referred to as the *DTBook file*. An audio-only book doesn't contain text and therefore will not include a DTBook file. The text content is marked up with a set of elements, similar to HTML, that indicate the information's structure and semantics. This markup is also known as *Daisy XML*. Daisy includes about 80 elements³ such as

- `<sent>` and `<p>` to indicate sentences and paragraphs;
- `<level1>` through `<level16>` to indicate structural components, such as chapters and sections;
- `` and `<imggroup>` to include images and image descriptions; and
- `<pagenum>`, `<noteref>/<note>`, and `<sidebar>` to mark page numbers, footnotes, and sidebars.

For the richest reading experience, the DTBook file should contain the full text, including footnotes and descriptions of all of the images in the book.

A Daisy book might optionally have images, stored in separate files. Daisy supports common image formats such as SVG, PNG, and JPEG. A Daisy book might contain recorded audio created by a person narrating or a synthetic speech voice. Daisy supports common audio formats such as MP3 and WAV.

When a Daisy book contains text, audio, and images, the synchronization file specifies the timing and order of presentation. The file conforms to SMIL 2.0 and is referred to as the *SMIL file* (pronounced “smile”). The body of the SMIL file consists of a series of `<seq>` and `<par>` elements. Children of a `<seq>` element are presented sequentially to the reader and children of a `<par>` are presented simultaneously (in parallel). In a Daisy book the children may be

- `<text>` elements that reference content in the DTBook file,
- `` elements that reference image files, or
- `<audio>` elements that reference content in the audio file.

For example, a `<par>` element might contain a `<text>` element that refers to a particular `<sent>` element in the DTBook file. The `<text>` element might be followed by an `<audio>` element that refers to a particular portion of an audio file by its starting and ending clock values. When the `<par>` element is played back, the player simultaneously displays the `<sent>` element of the DTBook file and plays the section of the audio file referenced in the `<audio>` element. If the `<text>` and `<audio>` elements are children of a `<seq>` instead of a `<par>`, then the `<sent>` element is displayed first, followed by playback of the audio file. Image playback is handled similarly by using the `` element. Figure 1a (next page) shows an example of a SMIL file controlling text and audio files. The `<seq>` and `<par>` elements can be nested in other `<seq>` and `<par>` elements to provide any required combination of sequential and parallel presentation of text, images, or audio.

The navigation control file, referred to as the *NCX file*, supports navigating to particular parts of the book. The NCX file summarizes the significant structures of the book. One part of the NCX information acts as a hierarchical table of contents. Another part of the file links to the book’s pages. This functionality permits the reader to browse the page list or jump directly to a given page. The NCX file might contain lists of illustrations, tables, equations, or other items of interest. The reader can

browse a given list and jump directly to an item. Items in a list don’t need to be of the same type; a list might be named *pendulum experiment*, for example, and refer to associated text, equations, figures, and tables.

The resource file specifies text, audio clips, or images for notifying the reader when encountering certain book features or structures. For example, the resource file makes it possible to have page numbers audibly announced if the user chooses to hear them.

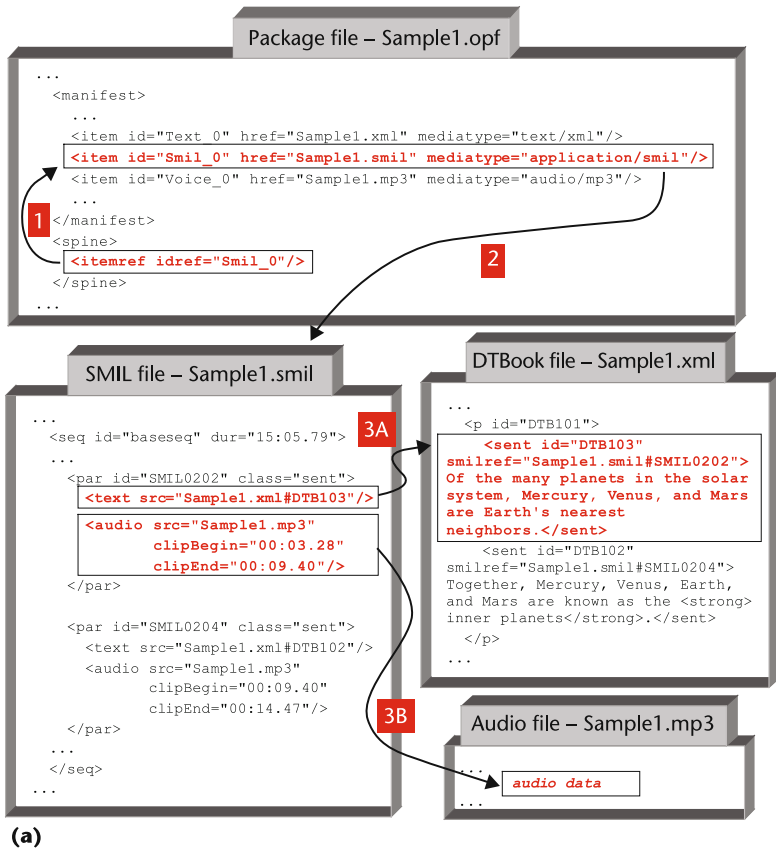
In addition to these features, Daisy books support formatting through both cascading stylesheets (CSS) and Extensible Style Sheet Language (XSL) styling, as noted in the presentation style file. The styling can be applied to the visual, audio, or braille presentation of the book’s content. Transform files are used when some information in a Daisy book includes a simplified fallback format for players that aren’t able to handle advanced content. Mathematics in the form of MathML, for example, has an image and text fallback format. The Daisy book can include XSLT to perform transformations into a simplified format.

The Daisy standard describes two other file types not included in the previous list and not part of the book itself. A media unit, such as a CD-ROM or flash drive, might contain more than one Daisy book. Conversely, a large book might reside on several media units. In these situations, the distribution information file describes how to locate the necessary information for each book. And for cases where readers want to set bookmarks at points of interest, they can do so with their own text or audio notes, which are referenced in a bookmark file and are sharable with other people.

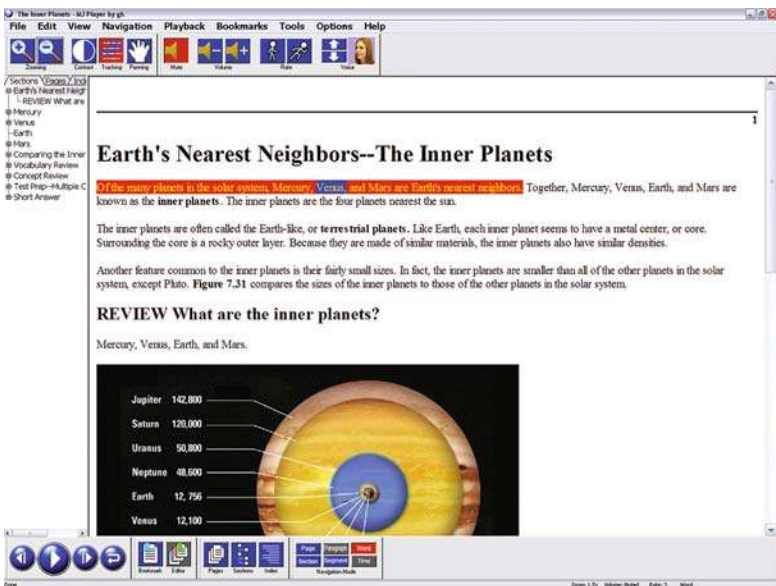
For reasons of performance or convenience, each type of information can be divided among several files. For example, the content of a large book may be split into several DTBook files. Also, when a book spans several CDs, there are further constraints on the files to minimize swapping between the CDs. Some information, like the package file, is duplicated on each CD.

Daisy book production

Typically, Daisy book production begins with an existing print book. There are different types of production processes: starting with text or starting with audio.



(a)



(b)

Figure 1. A Daisy player's (a) typical operation and (b) rendering of the content shown in (a).

For a text- and image-only book, the material must be available in electronic form. The publisher could supply the content as either typed or scanned pages. Next, the text is marked up with the DTBook element set.³ Some tools will automate this step; others will require human guidance. Images must be

converted to SVG, PNG, or JPEG and their textual description added if not present. Extensive and detailed mark-up produces a better reading experience but calls for more manual work. The *Daisy Structure Guidelines* provide directions and examples of recommended use of the DTBook elements.¹² Eventually, the mark-up produces a prototype DTBook file. A program reads this file to generate an actual DTBook file and all the associated Daisy files (package, SMIL, NCX, and so on). All associated data, such as navigation points, pages, footnotes, and so forth, are extracted from the marked-up text.

Usually, Daisy books produced in this manner don't contain audio. However, a text-to-speech engine can generate audio files during production to create a Daisy book with recorded synthetic audio. A different process model takes recorded audio, usually from a human reader, as the production master. Resulting Daisy audio-only books generally don't contain text or images. However, they do have some structure in the form of headings and pages, which enable navigation.

To produce a Daisy book with text and audio, it's necessary to produce both the marked-up text and the audio files, and provide synchronization between them. The production tool can link the audio and the text at the paragraph, sentence, or even at the word level. A SMIL file is generated that contains a link to each text unit and a link to the corresponding audio fragment. This way, text and audio in the resulting Daisy book can be rendered in parallel.

The range of production tools reflects these different production models. Tools can be simple, such as reading a Microsoft Word or similar document and using the styling information to mark up the text. Other tools are basically sound recording devices that directly produce a set of Daisy files. Some production tools allow for synchronization at production time of human recording and marked-up text. If audio and text were prepared separately, they can be joined afterward. Linking by hand is a labor-intensive job. Exploration of more automated ways of synchronizing text and audio shows promise.

Daisy has initiated development of a new set of production tools especially for the Daisy 3 format. The following tools are open-source

licensed under the Lesser General Public License (LGPL):¹³

- The Daisy Pipeline transforms documents between different formats such as Microsoft Word, Rich Text Format, XHTML/HTML, and older Daisy formats.
- The Urakawa multimedia authoring toolkit enables designing fully accessible content.
- The Obi Daisy production tool lets the individual Daisy reader (including print-disabled users) create Daisy content. This tool is based on Urakawa.

There are also several commercial production tools. These include

- Dolphin Publisher, EasyProducer, and Easy-Converter, which are production tools for consumers and professionals;
- PlexTalk Recording Software Pro, which annotates recorded audio with navigation links; and
- eclipseWriter, which is a production tool for professional and commercial use.

The number of Daisy-related production tools is growing rapidly. The Daisy Consortium lists several of them on its Web site (see <http://www.daisy.org>). As one example, Microsoft and Daisy have collaborated on a free add-in for Word that permits users to save a document as Daisy XML (a DTBook file).

Daisy book players

The reader's experience can vary widely depending on the type of book and playback system. There are two major categories of Daisy playback systems:

- hardware devices, which are usually small, highly mobile, audio-only systems with a limited user interface, much like a portable CD or MP3 player; and
- software applications, which are usually computer-based and have a sophisticated interface capable of showing text, images, and audio.

These categories are blurring as Daisy players are showing up on PDAs, mobile phones, and ultrasmall PCs. These are highly mobile, multiple-purpose devices capable of audio playback and text and image rendering. In addition, Daisy books can be converted to other formats for playback by iPods, Web browsers, and so on, or converted into paper braille. After conversion, however, the documents might lack some of the functionality that Daisy offers.

The Daisy Consortium maintains an up-to-date list of players on its Web site. In early 2008, the site listed approximately

- 20 hardware players, half of which support Daisy 3; and
- 20 software players, most of which support Daisy 3 and run on Windows, although a few run on Macintosh and Linux.

Some hardware players incorporate portable braille displays or PDA-like capabilities; most render books only with recorded audio. A few use synthetic text-to-speech to read text-only books. A couple are even capable of recording audio and directly producing a Daisy book from the recorded audio.

Typically, software players render the book text (if present) with some magnification, or use alternate colors to display the text. Most software players feature synchronized highlighting of text and audio. Some offer synchronized braille via a refreshable braille display. As one example of a software player, AMIS is a free open-source application maintained by the Daisy Consortium (see <http://www.daisy-for-all.org/?q=taxonomy/term/7>). It features easy localization to different languages.

To make the operation of Daisy players concrete, Figure 1a shows some of the files used by Daisy players and how they are interconnected. A player first opens the package file (top of figure) and finds the `<spine>` element (arrow 1). This element specifies the playback order of the SMIL file. The player uses the `<manifest>` element to locate and open the correct SMIL file (arrow 2). Linear playback basically consists of walking the SMIL file, playing back each child of the initial `<seq>` element.

The bulk of the SMIL file consists of a `<seq>` element whose child elements are `<seq>`s and/or `<par>`s. The children of a `<seq>` element are played back in sequence. When the first child finishes playback, the second begins, then the third, and so on sequentially through all the children. The children of a `<par>` element play back in parallel. All of its children commence playback simultaneously. The entire `<par>` element is finished when the longest running child finishes. In Figure 1a, the player is currently on the `<par>` element whose ID is SMIL0202 and contains a `<text>` and `<audio>` element.

The `<text>` element refers to a sentence in the DTBook file (arrow 3A in the figure). The DTBook file is opened and the element with ID DTB103 is located. In this example, the sentence “Of the many planets...” is found. The `<audio>` element in the SMIL file, meanwhile, refers to a segment of audio in an MP3 file (arrow 3B in the figure). Audio segment playback commences simultaneously with playback of the sentence in the DTBook file. Audio-only players ignore the `<text>` element and only use the `<audio>` element.

After playback of the `<par>` element, the player advances to the next child of the `<seq>`, which is the next `<par>` element shown.

Text playback of the sentence can take several forms: the sentence can be visually highlighted if a display is available, audibly rendered via text-to-speech, or translated and shown on a refreshable braille display. The audio segment can be played back more quickly or more slowly than the normal rate, or filtered for better recognition.

Figure 1b shows a typical rendering by a Daisy player of the files shown in Figure 1a. The book content is displayed in the main area and a navigation pane is available on the left. The navigation pane displays information such as section headings and page numbers and enables jumping to sections and pages. This information comes from the NCX file.

The current sentence is highlighted in one color and the currently sounding word is highlighted in another color. As each word is read, that word is highlighted so that the audio is always synchronized with the display. When reading advances, the next sentence is highlighted and the word-by-word highlighting continues. If a braille display is available,

the current sentence is displayed there, providing a multimodal user experience with sound, vision, and tactile renderings of the book.

Controls along the top and bottom of the screen allow the user to adjust display properties, such as magnification and color contrast, along with speech properties, such as volume and reading speed. These choices augment CSS information to produce the rendering shown on the screen.

MathML in Daisy

Historically, math expressions in digital book formats were typically displayed as images with pregenerated text. Using MathML in Daisy⁹ instead of embedded images provides many advantages, including

- magnification of an equation or changing its colors,
- tailoring the speech to individual needs,
- navigation and exploration of the mathematical structure,
- synchronized highlighting of math with audio, and
- translation to braille math code for use on a refreshable braille display.

Several files must be adjusted to include MathML in a Daisy book. The most interesting changes occur in the DTBook and SMIL files. Figure 2a shows a portion of these files. The `<text>` element in the SMIL file, which refers to the MathML (arrow 1), has a type attribute that signals MathML encoding in the DTBook file. The `altimg` and `alttext` attributes on the `<math>` element in the DTBook file are used for fallback behavior in case the player can't handle MathML natively. Existing audio-only players ignore the `<text>` element and use the `<audio>` element as their fallback behavior.

Figure 2b shows a rendering of text mixed with MathML. Here the user has increased the screen magnification to 600 percent and chosen a contrast color scheme. The math expression is rendered from MathML instead of displayed from a static image. This allows the player to smoothly draw the math at the

same size and with the same color scheme as the text. Direct use of MathML by the player also permits synchronized highlighting of math with audio, so that each part of the math expression is highlighted as it is read, as Figure 2b shows. In addition, the use of MathML permits readers to interactively explore the expression, piece-by-piece, via special keystrokes. None of these features are possible if the expression is rendered from a static image.

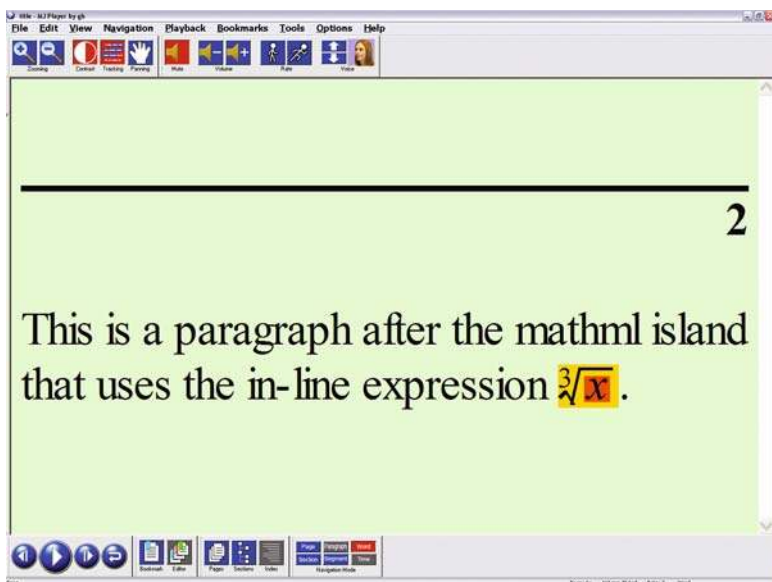
Although MathML is a recent extension to the Daisy standard, tools that work with MathML are already available. These include the MathFlow tools from Design Science for authoring and, for playback, the gh Player (used for the renderings in Figures 1b and 2b).

Daisy practice and experience

In most countries, Daisy is the main or only distribution format of books for people with visual impairments. They can listen to Daisy audio books and read Daisy textbooks using assistive technology such as text-to-speech, magnification, and braille display facilities. Synchronized audio and text highlighting benefits people with dyslexia, certain learning disabilities, and low literacy, as well as those learning a second language. People who use Daisy audio books include those with motor impairments ranging from arthritis to paralysis, as well as those who have temporary conditions such as an injury and those who are recovering from surgery.

Daisy books typically aren't available at the local library. Instead, most people with disabilities obtain them via a national library or nonprofit organization. Most books from these entities are for leisure reading. However, a growing number of educational and reference books are being converted to the Daisy format. In the US, the federal government mandate that publishers use the NIMAS format for K–12 educational material is dramatically increasing the availability of accessible textbooks.¹ The number of Daisy book titles available worldwide is well over 250,000 and is growing rapidly. Commercial sales of Daisy books began in 2005 and include one of the leading high school chemistry textbooks in the US (*Chemistry in the Community*, 5th ed., by the American Chemical Society).

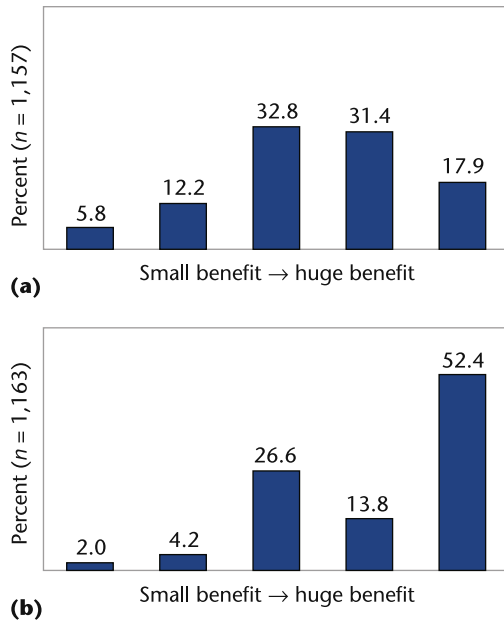
The readily accessible book structure of Daisy provides one of the major advantages



over other formats. The Daisy book format enables quick access to pages, subsections, footnotes, sidebars, and other structured information. Easy navigation is especially important for education and reference materials. An instructor, for example, might direct

Figure 2. (a) MathML references in the SMIL and DTBook files and (b) rendering of content from (a).


Figure 3. Perceived advantage of Daisy in comparison to audio CDs: (a) leisure books and (b) reference books.



students to a particular page number. Daisy provides this type of random access as well as access to specific words or phrases through custom searches. Moreover, reader-defined bookmarks permit personalized navigation points, which are useful in education settings and for reference books. In a recent diplom thesis,¹⁴ for example, Ludwig evaluated the features and possibilities of the Daisy standard and found that a significant number of the surveyed users perceived considerable advantages to using Daisy technology versus audio CD, particularly for reference books (see Figure 3).

A key advantage of the Daisy system is the diverse navigation possibilities, namely the ease with which the user can find his or her way around the text. The technology allows for easy navigation to find a particular section, passage, or even word, as a recent study by the Royal National Institute of Blind People (RNIB) found.¹⁵ One of the study participants remarked, "Excellent to have this tool, especially with full text Daisy books. All the bookmarking, easy navigation ... are excellent for

Figure 4. Rating of Daisy features on a scale ranging from: 0 (unimportant) to 4 (very important).



students," In support of this RNIB study, Ludwig found that bookmarks and navigation were what readers found most important, while the ability to control the speed of playback was not as important (see Figure 4).

In another study, the Royal New Zealand Foundation of the Blind did a study of Internet delivery of Daisy books aimed at acceptance by the elderly (ages 60 to 93).¹⁶ The participants found Internet delivery convenient; the players simple to use; and were positive, in general, about the experiment. One participant even exclaimed, "Oh, that's grand."

Daisy's future

The use of Daisy is expanding beyond books to include journals and newspapers. In January 2008, the American Physical Society and ViewPlus Technologies announced a pilot project to create APS journals in the Daisy format for use by scientists (see <http://www.viewplus.com/company/news/2061/>). As another example, the Pratsam project in Finland has developed an Internet-capable Daisy player that gives participants online access to nine daily newspapers (see <http://www.pratsam.com>). For these and similar projects, access to the Internet is important because distribution of new content by mail would not be timely. To facilitate Internet distribution, Daisy players and content providers must agree on a standard protocol. The Daisy Consortium Online Delivery Project is drafting such a standard (see http://www.daisy.org/projects/daisy-online-delivery/daisy-online_workarea). The International Federation of Library Associations and Institutions (IFLA) is studying the benefits of a global virtual library using Daisy.¹⁷

Daisy books have another benefit: the reader, not the publisher, chooses the book's format. Software-based Daisy players offer multimedia reading options, but a user might sometimes prefer an audio format and at other times might need a paper braille copy. The Daisy multimedia format allows users to choose which presentation format works best for them, from audio to braille to a mixed multimedia rendering. When coupled with online distribution from digital libraries, the Daisy standard holds great promise for giving people of all abilities timely access to a wide range of materials in a format that best suits their needs.

MM

References

1. S.E. Shaywitz and B.A. Shaywitz, "The Neurobiology of Reading and Dyslexia," *Focus on Basics*, vol. 5, no. A, 2001.
2. J. Elkind, "Using Computer-Based Readers to Improve Reading Comprehension of Students with Dyslexia," *Annals of Dyslexia*, vol. 43, 1993, pp. 238-59.
3. *Specifications for the Digital Talking Book*, NISO, 2005; <http://www.daisy.org/z3986/2005/Z3986-2005.html>.
4. R. Ausbrooks et al., *Mathematical Markup Language (MathML) Version 2.0 (Second Edition)*, World Wide Web Consortium (W3C) recommendation, Oct. 2003; <http://www.w3.org/TR/2003/REC-MathML2-20031021/>.
5. J. Ferraiolo, F. Jun, and D. Jackson, *Scalable Vector Graphics (SVG) 1.1 Specification*, W3C recommendation, Jan. 2003; <http://www.w3.org/TR/SVG11/>.
6. W. Raeder, "The Case for Braille," Nat'l Braille Press, http://www.nbp.org/ic/nbp/braille/case_for_braille.html.
7. The Eye Diseases Prevalence Research Group, "Prevalence of Age-Related Macular Degeneration in the United States," *Archives of Ophthalmology*, vol. 122, no. 4, 2004, pp. 564-572.
8. J. Ayars et al., *Synchronized Multimedia Integration Language (SMIL 2.0) - [Second Edition]*, W3C recommendation, Jan. 2005; <http://www.w3.org/TR/2005/REC-SMIL2-20050107/>.
9. N. Soiffer and K. Kahl, *Specification for a Digital Talking Book Modular Extension for Mathematics*, Daisy Consortium, 2007; <http://www.daisy.org/projects/mathml/mathml-in-daisy-spec.html>.
10. *Open Publication Structure (OPS) 2.0 v1.0*, Int'l Digital Publishing Forum, 2007; http://www.idpf.org/2007/ops/OPS_2.0_final_spec.html.
11. *NIMAS Technical Specification*, CAST, 2005; http://nimas.cast.org/about/proposal/spec-v1_1.html.
12. Daisy Consortium, *Structure Guidelines for Daisy 3*, 2008; <http://www.daisy.org/z3986/structure/SG-DAISY3/>.
13. Free Software Foundation, "GNU Lesser General Public License," 2007; <http://www.gnu.org/licenses/lgpl.html>.
14. K. Ludwig, "Das Hörbuch der Zukunft? Online-Studie zur Nutzerakzeptanz des Digitalen Hörbuchstandards Daisy [The Audio Book of the Future? Online Study of the User Acceptance of the Digital Audio Book Standard Daisy]," VDM Verlag Dr. Müller, 2007.
15. M. Cairns, "Case Study: RNIB and Dolphin Create Equal Opportunities in Education," 2005; http://wiki.euain.org/doku.php?id=wiki:case_studies:dolphin_and_the_davie_project.
16. R. Bijl, M. Schnackenberg, and C. Lansink, "Internet Delivery of Daisy Talking Books—Achievability and Acceptability for Blind and Vision Impaired Seniors: The New Zealand Experience," 2007; <http://www.rnzfb.org.nz/ourservices/readingandwriting/talkingbooks/DTB%20trial>.
17. E. Tank and C. Frederiksen, "The Daisy Standard: Entering the Global Virtual Library," *Library Trends*, vol. 55, no. 4, 2007, pp. 932-949.

Dennis Leas is the director of technology at gh, LLC, where he heads development of accessible software including the gh Player for Daisy. He represents the Daisy Consortium on the W3C Math working group and is a member of Daisy's MathML working group. Leas has a BS in electrical engineering from Purdue University. Contact him at dleas@gh-accessibility.com.

Emilia Persoon is a project manager in research and development at Dedicon, which is a founding member of the Daisy Consortium. Her research interests include IT-related subjects such as XML-based technologies, accessible multimedia exams, and production of Dutch math braille code. She is a member of Daisy's MathML working group. Persoon has an MS in computer science and in biology. Contact her at epersoon@dedicon.nl.

Neil Soiffer is a senior scientist at Design Science, where he focuses on math accessibility problems. He worked as a principal architect of MathML, developed Mathematica's WYSIWYG math editor, and chairs Daisy's MathML working group. Soiffer has a PhD in computer science from the University of California, Berkeley. Contact him at neils@dessci.com.

Michael Zacherle is a research associate and program supervisor at the Study Centre for the Visually Impaired Students at the Karlsruhe Institute of Technology. He is a member of Daisy's MathML working group. His research interests include design and evaluation of nonintrusive mobile and multimodal man-machine communication for blind users. Zacherle has a diplom in physics from the University of Karlsruhe. Contact him at michael@zacherle.de.