

Discussion Paper #1

6. Kay (2001 [1989]) explains Bruner's interpretation of Piaget as involving three components, each labeled one "mentality", or way of thinking about problems and situations. Each of these is reflected in the design principles of the Dynabook, as we shall see.

The first of the mentalities is the enactive component. Responsible for knowing where one is and manipulations of the environment, this component is reflected in the mouse. The mouse provides constant feedback (via a "pointer") and also allows the user to manipulate objects ("clicking", "dragging" and so forth) and so fits this role well.

The "iconic" mentality is the second component which recognizes, compares and configures. It is also said to be relatively concrete. In particular, it suggested to Kay to make use of pictorial representations. As the Chinese are said to say, "A picture is worth a thousand words", and Kay brought pictures to computer interfaces. These have the advantage of being easily apprehended as wholes (or perhaps as systems. See Bunge [1977] for a discussion of the characteristics and advantages of systemic as opposed to both holistic and reductionistic or atomic approaches. The gist is that the systemic approach has the virtues of both reductionism and holism.) that purely textual interfaces do not. From here we get windows, icons and other by now familiar items of graphical user interfaces. A special case of this mentality is the importance of comparison, which lead to the innovation of multiple windows. Furthermore, in order to facilitate the resulting movement between windows for comparison and interaction, Kay made the decision to design a *modeless* interface. This involves a bit of the first feature. In this style of interface the user is always "in compare mode" and doesn't have to worry about the restrictions on such found in other systems - no more "press

escape to edit" as found in The Bank Street Writer¹. The essential feature of the modeless interface is the ability of the user to take any action with the software at any time. (This is an ideal type, of course.)

The third mentality, "symbolic" is characterized as involving the tying together of long chains of reasoning and is (relative to the "iconic" mentality, one assumes) more abstract. It was here the Smalltalk programming language was to play a role. One of the earliest examples of an object oriented programming language (Downing and Covington 1992), Smalltalk was well suited to embodying the symbol ideal. As with all programming languages, it permitted the user to express her own chains of reasoning to solve particular problems. However, Smalltalk, because of it being object-oriented, allowed a curious tying together of all three mentalities. An object oriented programming language is also good from the enactive perspective as it allows data structures to perform their own operations, "knowing" what they can and cannot do. This can save the headache of purely procedural languages which require (as in C, for instance) the programmer to keep track of what structures make sense to pass to which functions and the nightmare of "pointer types". Object oriented programming languages also speak to the iconic mentality as well, as it is held that a graphical user interface can be implemented and maintained with greater ease in an object oriented language. Rather than have procedures that are invoked more or less sequentially, one instead can have an "event driven" approach" where each interface widget is coded by an object that sends and receives messages. Procedural language hacks to support event driven programming were later to be developed, but they are not as elegant. For instance, as illustrated in the *Macintosh C Programming Primer* (Mark and Reed 1992), one needs to use an infinite loop which passively waits for messages to be put on an event queue that the program can then read.

¹ This example is strictly speaking an anachronism. The Bank Street Writer (Brøderbund Software, 1982) was released after Kay's work at PARC. This is an excellent illustration how long it takes sometimes for technical innovation to reach the marketplace.-

By adopting an object oriented language from day one, the design of the Dynabook would be able to avoid these unfortunate hacks.

This tying together (though not just as pertains to Smalltalk) was expressed by Kay in a slogan: "Doing with Images makes Symbols." (2001 [1992], pp. 128) (capitalization in original.) Thus many of Kay's design decisions were not isolated, but connected systemically, yielding an important "structure of design" for the proposed Dynabook.

7. Vannevar Bush (2001 [1945]) discusses seven key features of how humans think and their role in the development of the memex model. The first of these concerns humanity's prodigious output of information. The memex thus must be able to handle this output. Humans also not only keep this information to themselves, but also share it with others Bush was dissatisfied with our ways of transmitting and reviewing information. This leads to the suggestions of using photocells to scan information and relay combinations for analyzing it. These draw upon our own capacities as visual and analytical creatures. The discussion that follows deals both with the memex proper and some of its more general presuppositions and parts that are found in existing devices. I take Bush as showing how these "fit in" to the later proposed memex model.

The human animal also continually extends thoughts; she must be able to consult the stored record mentioned above. Bush continues this train of thought to discuss possible developments in photography, in particular miniaturization and use of film with more capacity and other details. Transmission of images is also vital for the sharing we have previously mentioned. Thus Bush's proposal includes facsimile transmission. This Bush suggests shall be merged with television to form a certain form of "dry" photography.

In order to meet storage needs for these vast quantities of information, the memex will use microphotography of an extreme kind. This will not only render human knowledge more compact, but also cheaper (particularly in transportation costs). Humans are

creatures who often think in cost-benefit terms, so this bonus that comes along with storage concerns is also related to the way they think.

Humans often think aloud, or need records to share or consult later. Bush thus suggests that the use of a speech recognition system such as one demonstrated at a then recent world's fair would meet the demands of part of human thought. Similarly for repetitive thoughts; a submachine or component to handle some of these, such adding columns of figures would thus also help immensely. Bush even suggests that we adopt a notation to simplify the input of such into the machine, suggesting perhaps the punch card might be a way to go here. These calculator-like components are to be able to perform a wide range of operations that humans can, and also, like humans, possess a memory and a flexibility of doing various sequences of calculations (including advanced mathematics for scientific or technological research) with the same basic structure. Here Bush anticipates the computer program, suggesting that the pattern of calculations to be performed by his proposed device be determinable by considerations external to the structure of the machine.

Humans are also logical (at least some of the time) and so the machine must reflect that as well, according to Bush. For this reason, it would be useful to have the machine capable of manipulating logical ideas and symbols. But humans are also not just fact accumulators and referencers, as much of the above considerations have been concerned with. We also are fact organizers. Here we have an explicit recognition by Bush about his proposal making use of a psychological hypothesis² - that the human mind works by association, rather than by subclassing or "can-only-find-it-in-one-place" library like classifications. It is in this context that the memex is explicitly discussed.

He recapitulates the importance of microfilm and dry photography we have discussed earlier, and mentions that the memex does work

² In the previous cases it seems that the psychological hypotheses and views that ground each part of Bush's proposal are not as explicitly held or discussed by him.

by conventional indexing in addition to the associational "indexing" he is about to discuss the implementation of. One of the other features of the memex that is also "psychological" in character concerns the multiple rates in which it can display information. Sometimes humans think over things fast, other times they want to examine in greater detail. To this end, Bush suggests that the memex have the ability to replay its contents at various speeds.

After this discussion, we finally get to the *pièce de resistance* of the memex, the ability of it to tie two items together at will. Bush himself calls this the "essential feature" of the memex (2001 [1945], pp. 150). Each tie is labeled with a code word and then hereafter when ever either item in a tie is being viewed a button on the memex' keyboard will allow transfer of focus to the other item. This allows the user of the device to draw up great associational networks between items, to see what other people find in the way of associations (as when consulting a prepackaged encyclopedia, for instance) and to also take advantage of the memex's memory abilities to make sure that the trail created does not fade.

Bush also speculates on other input and output mechanisms to the memex, noting that we are learning great new things about the transmission of electrical signals, that our nervous system (our very thought!) works electrically and hence it might be possible to input and output to the brain directly.

Except for the last item, which remains in development stages at best, virtually all of these features are found in modern computer technology. The vast quantities of informational storage is certainly true. The consumer can buy a random access storage device that stores 320 gigabytes of information as of this writing (Shareholder.com 2002). Being random access, this hard disk not only meets the latter feature of Bush's proposals but the retrieval and even the transmission angle. The latter comes in with the size of storage issue mentioned earlier; it is certainly easier (and cheaper) to send a hard disk than the books it might contain. Retrieval is straightforward from a hard disk, though it

is not self contained as requires cables and a general computer system (with the appropriate software and other hardware) to make use of it. This hard disk also fulfills the cost-benefit consideration previously discussed.

A modern computer system certainly has much in the way of information transmission capability. Many different networking hardware standards exist, with different protocols supported over these technologies (e.g. AppleTalk or TCP/IP, both over Ethernet).

However, it is only by extended analogy that we have access to Bush's proposal of microphotography. It is not true that we store our information in a photographic way, per se. This paper is not stored as if it were a photograph on my computer's hard disk. Curiously, however, the modern computer in a certain respect reverses Bush's proposal: we now have digital cameras and scanned photographs that store photos in "computer formats" rather than computers using general photographic formats. Bush was right to anticipate a blurring of "data storage methods" in this sense, then.

Bush also proposes the use of speech recognition for various purposes. This has been available for (at least some) personal computers as a vendor shipped feature since 1993 with the introduction of Apple's (then) Centris 660av and Quadra 840av. As an add on, it dates prior to this in some limited use applications for the Commodore 64 6 or so years earlier. Nevertheless, the ease of dictation and general use of such systems is not quite at the point that Bush would have hoped. Aftermarket speech recognition software (such as IBM's ViaVoice) is more sophisticated in certain respects than much vendor supplied speech recognition, but is still not ready for perfect dictation uses yet. That said, we are progressing to Bush's ideal, and further, in a slightly more sophisticated way than he suggested. His use of the Voder involved transcription into a special phonetic alphabet much as was used by existing human stenographers. ViaVoice and other similar applications do not require the user to further translate the output - except of course in the case of errors in recognition.

Needless to say there are many calculation programs available of remarkable sophistication. Scientific calculations of all sorts capable of computer resolution. Furthermore, something Bush did not directly dream of that is of equal importance in relieving tedium in scientific and technological contexts is symbolic manipulation. Mathematica is regarded as one of the best products in this area and is also affordable for individuals, not just large institutions or governments. Similar considerations apply to logic proper; logical manipulation software is also available not only for desktop computers (e.g., Bertrand) but available as a server application to be distributed to clients (the CMU Proof Tutor, for instance).

Bush's desire to have the machine be able to display information at various speeds is also obviously realized by our contemporary computers. In fact, our computers often display information too fast, which sometimes leads to design decisions to slow them down (for instance, the "scroll bar" throttle in MacOS is slowed down, unlike the one in most versions of Windows). Nevertheless, the scrollbar meets this Bushian ideal, as do settings for "slide show" programs like PowerPoint.

It is only relatively recently that we have seen the ubiquity of something similar to Bush's informational ties. The most common of these is the hyperlink, first made relatively popular in Apple's HyperCard (first released in 1987) and then made ubiquitous by the World Wide Web and its Hypertext Transfer Protocol and Hypertext Markup Language. It is important to stress that most hyperlinks in these currently popular implementations are unidirectional; I click on a link on a web page and cannot usually go back to where I was by the same method. Of course, there are current technologies that do allow multidirectional links (as Bush suggested), such as some XML approaches (Williamson 2002), but these are not nearly as prevalent as the other technologies at the time of writing. Thus, another one of Bush's dreams remains only partially fulfilled.

8. Nelson's notion of "hypertext" (2001 [1974]) has two key

features, encapsulated in its basic definition: hypertext is non-sequential writing. Since structures of ideas are not sequential and tie together in arbitrary ways, hypertext is designed to reflect the structure of ideas (said to be nonsequential), rather than the sequential nature of speeches and (most) books³ .

From this basic definition comes several more nuanced notions. Nelson introduces "basic", or "chunk style" hypertext, "collateral hypertext", and "stretchtext". These are styles of hypertext: basic hypertext amounts to a footnoting or referencing scheme; collateral hypertext involves annotations and other texts. Stretchtext is in turn characterized by continuous change. After introducing these styles, Nelson then discusses species of hypertext documents composable in these styles. Very broadly, "fresh hypertexts" (i.e. new ones written for a specific purpose) combine to form "anthological hypertexts", which are analogous to anthologies in conventional publishing. These in turn combine to form the "grand hypertext", which is the (mereological) sum, with appropriate additional interconnections, of all the anthological hypertexts on a given subject. It is not clear whether there could be more than one of these per subject; perhaps so, given that one can connect ("link") the anthological hypertexts in an arbitrary number of ways.

Once this is all set up, with people having access to many grand hypertexts, the question of "who will use these systems?" comes to the forefront. It is here that Nelson first compares hypertext to the processes foreseen by Vannevar Bush. Both Nelson and Bush feel that the community of scholars has become too disjoint, both internally and to human communities at large. Thus it is only fitting that everyone be able "to put them in the service of truth and learning", as Nelson says. (2001 [1974], pp. 161)

³ There are nonsequential books, such as the *Choose Your Own Adventure* series of novels and even stochastic books, like the *Fighting Fantasy Gamebook* series. The epitome of the latter is "Creature of Havoc", by Steve Jackson (Jackson 1986) as the "winning resolution" of this book necessarily requires several "successful" dice rolls. The present author is not familiar with any non-fiction explicitly nonsequential books, but is actually in the (slow) process of writing one with a friend.

From here Nelson moves on to comparing his proposals to the work of Douglas Engelbart. Both Engelbart's mouse and his NLS system are points of comparison with Nelson's hypertext. Hypertext involves making use of connections between the representations of ideas. Thus in order to utilize the connections somehow, a user of the hypertext system must be able to "activate" these connections in some way. Nelson's proposal for hypertext entails this necessity but does not give many details for how this is to be accomplished. The mouse is thus complementary to the Nelsonian hypertext system.

On the other hand, the NLS system is an example of a hypertext system in two ways. First, it involves putting a structure beyond the usual beginning-middle-end on text. The hierarchies that NLS creates add this further structure. Second, since the hierarchy can be display in various arrangements, the hierarchy itself is subject to internal reorganization. This makes it similar to the hypertext approach, where the internal structure of a "document" (broadly construed) is subject to (in some sense, created by) the user following connections at her whim and interest.

Nevertheless, there are a few dissimilarities between the two approaches that Nelson draws the reader's attention to. Nelson points out that NLS does default to imposing a lot more structure than he would like. This setting is said to be changeable, but the hypertext system would default the other way around - let the user decide to impose more structure rather than force her to use more and require her to change less if desired.

9. Berners-Lee's (2001 [1989]) prototype proposal for the "world wide web" is for a system that grows and changes as the information it structures grows and changes. He summarizes this in the maxim: "the method of storage must not place its own restraints on the information." He likens this approach to a diagramming approach, where relationships between things are represented with circles and arrows. In the "web", a circle becomes a node and the arrows links. Each node can be a note, an article or a comment, and can be textual or graphical. The link

then can represent one of several key relationships. Berners-Lee lists five: is a part of, made, refers to, uses, and is an example of. These are likened to circles and arrows on conventional ("paper") diagrams such as dataflow diagrams or organizational charts. Finally, the system must allow the user to find information without knowing what she is looking for specifically.

Berners-Lee's approach is contrasted with trees and keywords, which both have their limitations. The solution to all these considerations is hypertext. This will be used to put "hot spots" in documents, both static and temporary. This is more a function of what is referred to than the link itself.

Furthermore, the "web" must not only have the structural features discussed above, but must also have six more functional requirements. Berners-Lee suggests that remote access across networks is necessary, that different systems be able to access the same data, that systems be linkable together without central coordination, that existing data be still accessible, that users can add their own links to and from public information as well as annotate links and nodes privately. Finally, while the proposal at minimum concerns ASCII 24x80 terminals, graphics would be a useful optional extra.

Superficially, this proposal does resemble the contemporary World Wide Web. For instance, the current World Wide Web does make use of a node-and-link structure, broadly construed. However, the distinction between static and temporary links is almost nonexistent. That is, temporary links just lead to page not found pages, or similar error conditions, rather than the link itself disappearing after a certain time. Structurally speaking also, "very" dynamic content on the current Web is a profusion of standards and approaches. Some of the approaches include using other protocols than HTTP itself (such as the RTSP for Quicktime Streaming Server or the embedding of Java applets). Others include having serverside software "manufacture" content on the fly by parsing configuration files as is the case with Active Server Pages. Scripting languages (such as Javascript) embedded within the HTML served by HTTP are also another solution.

As for the functional requirements, let us look at each in turn. The remote access feature is well established; one can even access, in principle, the World Wide Web, from war torn and desperately poor Afghanistan⁴. Internet itself has incorporated many other networks (e.g., FIDONet), however, many of these do not have access to the Web through their subnetwork. On the other hand, corporate, personal and educational internets can easily join Internet and gain access to the web through many service providers of various scopes and pricings. The non-centralization aspect of Internet is inherited by the Web. There are several organizations, however, that provide some centralization of some aspects of the network, such as ICANN (ICANN 2002) and Internic (Internic 2002). It is possible to use the web without paying much attention to ICANN and such bodies, however at some level one's ISP or the like must on pain of anarchy. In this way there is no enforced central control, but instead there is a voluntarily imposed one. It is vital to realize, however, that organizations such ICANN are themselves distributed, so they are only functionally centralized. Finally, it is important to realize that there is increasing central control over access to parts of the World Wide Web, despite great freedom to connect to it. This is especially true in places like China, which has banned popular search engines like Altavista. However, there are rumblings of moves in this direction elsewhere; the (struck down) Communications Decency Act in the United States made Internet Service Providers more susceptible to central control of control.

Access to the existing (pre-Web) data was not maintained by HTTP itself. Earlier methods of access, such as Gopher and FTP are linkable to HTML documents and hence accessible "on the web" in certain sense. However, these earlier protocols are instead handled either by the Web browser understanding the protocols or by handing off the task of dealing with them to a "helper application." In the case of more specific databases and

⁴ My cousin Vance was in Afghanistan with a Canadian relief organization for some time, and was able to access a mail server via satellite connection while there. This access would, in principle, allow use of the Web, though in practice would be horrendously slow.

resources, ones that did not use well established protocols, there was little backward compatibility at all. The McGill University library system for many years was such; it was hosted on an MVS system and users would tn3270 to access the card catalogue software. This was not directly integratable with the Web and only more recently was the MVS system discontinued and a more up to date, Web-usable system was developed.

Private links are handled a bit differently than perhaps Berners-Lee had in mind. He writes (2001 [1989], pp. 198-199):

"One must be able to add one's own private links to and from public information. One must also be able to annotate links, as well as nodes, privately."

One can certainly create links to public information; one is free to create a page containing said links. However, hosting that page very often requires using someone else's server (e.g. that of a commercial ISP); so "one's own" is not quite correct. If one uses someone else's server, the question of privacy is then raised; one can use authentication tools and what not to help insure privacy, but no such tools are perfectly "hacker proof." The web browser itself has taken over this one of Berners-Lee's functions in some respect, allowing the user to make a "bookmark" to a desired page. Once again, however, the browser is in some ways outside the system and as such this approach does not quite capture what may have been intended. Bookmarks (like personal pages) allow annotation of links; one can name the link anything one wishes. However, it is harder to annotate nodes from a bookmark; in practice people simply annotate the node that results from following a bookmark by instead annotating the link. A personal Web page does allow annotation of both link (via the link name field in HTML) and the node (by surrounding text), but suffers from the limitations of personal pages in other respects discussed above.

Berners-Lee's last functional feature was not terribly prescient. Some Web pages these days work fine (so he was not totally mistaken) in "text only" browsers such as Lynx, but they appear to be a vanishing minority. More critically, Berners-Lee failed to realize how popular "bells and whistles" would become. There are

now hundreds of plugins for popular browsers (see, e.g., Netscape 2002), which integrate with HTTP and provide the appearance of the protocol supporting full motion video, animation, synthesized and recorded music, etc. Further, even the HTML standard has grown to include many non-ASCII features: text styles like bold and italic; non-ASCII characters such as the bullet, accented and ligatured characters, text colouration, and much else.

10. Morton Heilig's "cinema of the future" (2001 [1955]) anticipates developments in multimedia and virtual reality in eight ways, some of which do so more than others.

First, both make much use of large visual stimuli. The typical cinema screen (according to Heilig) spans 5% of one's visual field; his proposal, and that of VR, is to use 100% of the visual field. This, of course, is not realized in most multimedia applications, but is in some and that's sufficient for at least a weak anticipation.

Second, Heilig suggests that three dimensionality be incorporated. This is again found primarily in VR approaches; using existing software a chemist can reorient a molecular model in space in front of her and examine it from various sides. Simulations of 3D in 2D planes (as on computer monitors) is, however, very prevalent in multimedia contexts. Literally hundreds of games make use of this; "serious" applications such as travel software and surgical trainers use technologies like Apple's QuickTime VR for producing 3D effects out of 2 dimensional photos.

Colour and movement are also amongst Heilig's proposals. These are found in virtually everything that could be called "multimedia" at the present time. Importantly, colour has not always been present in such - many early "multimedia" applications ran on black and white MacOS machines with the use of HyperCard. Movement comes in several flavours, such as changes of points of view, common in "replays" during computer games, or changes in the position of an egocentric participant, which is common in many sorts of so-called "3D" games. Multimedia applications also allow views from impossible points of view; views from where no agent could be -

for instance, above a unit of troops as is common in a Warcraft game.

Heilig then turns to sound for his fourth group of suggestions. His recommendations of music, rhythm, words, etc. are also well established parts of contemporary multimedia - both synthesized (Macintalk, MIDI) and recorded (MP3 audio, etc.) sounds are prevalent. The latter is currently in ascendance; a perusal of the games in the Info-Mac hyperarchive (Info-Mac 2002) show an interesting trend starting with sampled sound, moving to synthesized ones (MIDI and SoundTracker formats borrowed from the Amiga) and then back to sampled sounds, this time with highly efficient compression that maintain high quality playback (MP3). Needless to say, all of these can be used in the more "virtual reality" like contexts.

The remaining three senses are much less developed. Research (for instance, at the University of British Columbia, in Vancouver, Canada (UBC Robotics and Control Laboratory 2002) on haptic interfaces continues. Some of the research in this general area has already produced fruits in the context of force-feed-back joysticks and pressure sensitive Wacom tablets. Nevertheless, more detailed stimulations of pressure and texture sensors in the skin is rather a long ways of.

Further, taste and smell are almost never a concern in contemporary multimedia applications. This is because these senses are chemical in nature, and barring direct stimulation of nerves, it does not seem likely that many multimedia products will wish to incorporate them for this reason. Nevertheless, there are a few exceptions, particularly expensive rides at theme parks and such. Heilig's call for the air to be filled with odors has not reached common implementation.

A more general consideration that Heilig makes is of not limiting the user to mundane environments. He suggests that new impressions such as those picked up by supersensitive microphones, be used to create environments that no human has experienced. This too is realized both in games of all kinds from Arkanoid to a VR

Mechwarrior game, and in more serious applications such as surgery training, anatomy tutorials and chemical modeling software.

Bibliography

Berners-Lee, Tim. 2001 (1989). "Information Management: A Proposal". Reprinted in Packer and Jordan 2001.

Bunge, Mario. 1977. *Treatise on Basic Philosophy*, Vol. 3: *The Furniture of the World*. Dordrecht: Reidel.

Bush, Vannevar. 2001 (1945). "As We May Think". Reprinted in Packer and Jordan 2001.

Downing, Douglas and Covington, Michael. 1992. *Dictionary of Computer Terms* (3e). Hauppauge: Barron's Educational Series.

Engelbart, Douglas. 2001 (1962). "Augmenting Human Intellect: A Conceptual Framework". Reprinted in Packer and Jordan 2001.

Heilig, Morton. 2001. (1955). "The Cinema of the Future". Translation by Uri Feldman in Packer and Jordan 2001.

ICANN. 2002. "About ICANN"
<http://www.icann.org/general/abouticann.htm>.

Info-Mac. 2002. "MIT LCS's INFO-MAC HyperArchive".
<http://hyperarchive.lcs.mit.edu/>.

InterNIC. 2002. "InterNIC Frequently Asked Questions (FAQS)".
<http://www.internic.org/faqs/index.html>.

Jackson, Steve. 1986. "Creature of Havoc." Volume 24 of the *Fighting Fantasy Gamebook* series. London: Penguin Books Ltd.

Kay, Alan. 2001 (1989). "User Interface: A Personal View". Reprinted in Packer and Jordan 2001.

Mark, David and Reed, Cartwright. 1992. *Macintosh C Programming Primer: Inside the Toolbox Using THINK C*. Vol. 1, 2e. Reading: Addison-Wesley.

Nelson, Ted. 2001. (1974). *Computer Lib/Dream Machines*. Excerpt reprinted in Packer and Jordan 2001.

Netscape. 2002. "Browser Plugins".
<http://wp.netscape.com/plugins/index.html?cp=brictrpr3>.

Packer, Randall and Jordan, Ken (eds). 2001. *Multimedia: From Wagner to Virtual Reality*. New York: W. W. Norton and Co.

Shareholder.com. 2001. "Maxtor Driving Capacity-Centric Enterprise Apps With "Super-Sized" ATA Drives"
http://www.shareholder.com/maxtor/ReleaseDetail.cfm?ReleaseID=89588&reltype=Product&maxtor_section=press.

UBC Robotics and Control Laboratory. 2002. "Robotics and Control Laboratory." <http://rcl0.rcl.ece.ubc.ca/>.

Williamson, Heather A. *XML From A to Z: A Quick Reference of More Than 300 XML Tasks, Terms and Tricks*. Redmond: Redmond Technology Press.