A 3–5 Year Computing Technology Forecast For The Academy

Technology Subcommittee
Advisory Committee on Academic Information Technologies
Portland State University
Portland, OR, USA 97207–0751

Subcommittee Members:
Judy Anderson <andersonj@pdx.edu>
Mark Gregory <gregorym@pdx.edu>
Bart Massey <bart@cs.pdx.edu>

January 21, 2002

Abstract

This document is an attempt to identify some of the key computing technology changes possible over the next 3–5 years, with an assessment of their likelihood and impacts. An emphasis is placed on those factors unique to the Portland State University environment.

Academic institutions are beginning to provide service for a new generation of users. These newer users expect institutional support for innovative new computing, communications, and interaction technologies in their homes and offices. At the same time, policy concerns such as the Digital Divide, environmental impacts, and issues of large-scale computing place new pressures on the institutional infrastructure. Finally, the technical and legal issues surrounding intellectual property and data interchange are fundamentally altering the computing infrastructure.

As with any such document, especially in the fast-moving world of technology, these sorts of prognostications are inherently imprecise. Nonetheless, there is little doubt that the academic world will continue to be pervasively altered by the introduction of new technologies. Any information or insight that can lead to productive outcomes from this process may be of major value.

1 Introduction

The process of identifying technology impacts on the academy can be fruitfully subdivided into two parts: identifying relevant technological factors, and understanding their impact. In this document, the impact of several key factors are discussed: hardware and infrastructure (Section 2), user interaction (Section 3), and software (Section 4). These are then summarized and some conclusions are drawn in Section 5.

There is an enormous amount of technological change in store for the computing world in the next few years, and a potentially huge set of impacts on the academic environment. Those factors presented here were selected for several reasons. First, they are reasonably likely to happen, at least within the 5-year time frame. Second, they are quite likely to have large impacts upon the academic situation. Finally, many of them are not currently the subject of major attention in the mainstream media: a few of the changes presented here may be surprising to even a technologically knowledgeable reader.

2 Hardware and Infrastructure

The easiest factors to predict, and in some ways the driving factors for the rest, are those in the arena of computing hardware and infrastructure. It is easy to forget, in the age of Virtual Everything, that computing is built on computers. The next few sections discuss some of the key ideas in the arena of computers, networks, and support structures.

2.1 Moore's Law and the Desktop

Moore's Law [11] is due to electronics pioneer and Intel founder Gordon Moore, who suggested in 1965 that the computational power available at a particular price doubles every 18 months. This extrapolation of then-available data has proven astoundingly accurate over the succeeding 35 years [14]. In part, this accuracy is because the industry has made a conscious effort to maintain this growth rate: software developers use Moore's Law to predict what computing power their software will be able to use when completed. Hardware developers use Moore's Law to gauge where their competitors will be in the future, and thus where they must be to compete.

According to Moore's Law, a $500 PC in the year 2004 will sustain 2 to 3 billion instructions per second, and will have 512 megabytes or more of memory. Memory speeds grow much more slowly than CPU speeds with time [10], so main memory performance will likely be the limiting factor in overall performance.

There are several obvious consequences of this increased power, and a few not-so-obvious ones. First, every machine in an academic setting needs to turn over every 3–5 years: machines older than this are hopelessly outdated. Because software developers count on the availability of higher performance, even upgrades of existing applications that add little or no new functionality may require substantially more computational power. The oft-repeated mantra of not upgrading is no solution to this problem: the extreme labor costs of supporting a heterogenous environment of unreliable and difficult-to-maintain hardware and software swamps the savings due to decreased hardware acquisition.

Indeed, one trend already appearing is the division of tasks into two classes: those that are not compute bound and never will be again, and those that can still use more computing power. Common tasks like word processing and spreadsheets have largely fallen into the first category. However, it is reasonable to expect that new demands for improved user interface technologies (Sections 3.2 and 3.4) and an increasing level of application "intelligence" (Section 3.3) will push these problems back into the second class at least in the short term. Eventually, the only problems remaining in the second class are likely to be those of simulation, modeling, and analysis (Section 2.4): this is unlikely to occur within the 5 year time frame of this document, however.

An odd result of Moore's Law is that there is a subtle tradeoff between purchasing machines on which work can be done now and waiting for "better" machines to be available. Preliminary analysis [6] indicates that the temptation to wait should generally be avoided: the immediate use of the machines is more important than the eventual savings.

2.2 The Digital Convergence

Up until the dot-com implosion [4], it was a commonplace assumption that televisions, cell phones, desktop computers, Personal Digital Assistants (PDAs), and the like would become subsumed by "information appliances" performing all of these functions and more in a single, integrated device. This still appears to be true, but the time scale is very difficult to predict.

2.2.1 The Handheld Explosion

The slow integration of handheld functionality—PDAs, cell phones, portable media devices such as music players and cameras, *etc.*—is already visible today, and will be only accelerated by the consequences of Moore's Law and the increasing sophistication of the marketplace. The consequences here are potentially enormous: many students are likely within 3 years to own and regularly use a portable device that provides all of the above-listed functionality, as well as reasonable-bandwidth wireless Internet access (Section 2.3.1).

One element that is difficult to predict and difficult to deal with is the "Digital Divide" [15] that may be especially visible in this area. If the most functional devices are cheap enough, almost every student will be able to afford one, and their use in classroom settings may be easily institutionalized. Unfortunately, the current class of highly functional devices appears to be in about the $300–500 range, which has greatly limited their student availability. An instructive and encouraging example is the pocket calculator: these devices were initially quite expensive, but rapidly became incredibly cheap and therefore nearly ubiquitous. In the near term, unfortunately, the problem of competition between students carrying metaphorical slide rules, calculators, or neither is likely to arise, and difficult to deal with.

2.2.2 Better Homes and Offices

The other front for the digital convergence is that of "fixed" computing environments such as homes and offices. The increased power of home computers discussed in Section 2.1 together with the increase in support for home broadband discussed in Section 2.3.2 means that more types of work-at-home options will become available to teachers, students, and staff. Notable impacts of this change include greatly increased demands on reliability and availability for campus network and computing infrastructure: the effect of a single-point problem or failure on the home or mobile user may be substantial.

2.3 Networking and Communication Technologies

Perhaps one of the biggest and most obvious changes in computing hardware and infrastructure over the last 10–15 years is the increasing quality, diversity, and ubiquity of computer-to-computer communications. Computer networking today is everywhere, in every media, and there is no sign of slowing in the near future.

2.3.1 The Wireless Revolution

The prediction of imminent large-scale wireless networking has been made repeatedly over many years. The current predictions, however, are supported by not one, but two actual technologies currently in production use. These technologies are IEEE 802.11 [7] and Bluetooth [2], and they differ from their predecessors in several important ways. They are true general-purpose packet-switched networking architectures designed from the ground up to be integrated into the Internet. They are supported by compatible (or at least mostly-compatible) products from a broad array of vendors. They are inexpensive and simple enough to overcome barriers to entry. Finally, their range and bandwidth are sufficient to support a reasonable range of applications. This said, the technologies are not equivalent. While 802.11 is suited for a wide variety of uses, Bluetooth is a less mature technology intended primarily for short-range, low-power, low-cost applications.

In addition, the profusion of handhelds, laptops, and other mobile and portable devices makes the demand for reasonable wireless computing solutions intense. Users would like to have the same kind of computing experiences on their portable device that they have on their fixed one, and to a large extent seem willing to pay for this desire in terms of price, complexity, and limited reliability. For all these reasons, it is likely that wireless computing will become widespread within the 3–5 year time frame. Indeed, Portland State University already has a substantial campus 802.11 infrastructure in place. Although the number of users is currently limited, this is already starting to change: as this document is being written, major new installations of 802.11 are appearing on campus, with more than 50% of the campus area already covered.

The impacts of this change are large. For example, the type and amount of infrastructure costs for campus networking will change significantly. Vast sums are currently spent to install and maintain the wireless plant, but this infrastructure does not generally require sophisticated maintenance above the physical level. As the demand for wireless increases, cost of installation and physical maintenance will likely be overcome by cost of maintenance of the software and management infrastructure needed to coordinate, secure, and serve a large number of novice wireless users.

Another obvious impact is on the pervasiveness of computing: wireless is a key component enabling full-scale computing in environments, such as classroom seats, previously deemed impossible. It may be difficult to predict the most important changes in this arena: the course taken is likely to depend more on social goals and ideals than on technology limits.

2.3.2 More and Faster Networking

To keep pace with the growth of computing power and the increasing reliance of society on digital networks, the capacity, performance, and availability of the network infrastructure must continue to rapidly increase. In addition, a form of feedback much like that common in the microprocessor world also occurs in networking: network capability grows to support cutting-edge applications, leading to a new generation of cutting-edge applications designed to utilize fully these new network capabilities. In the current generation, these applications are centered around video and audio, for example the coming replacement of much circuit-switched telephony with Internet-based communications.

In addition, the increased adoption and use of networking has led to huge increases in the amount of hardware devoted to linking computers: the quantity and speed of network switches on the PSU campus, for example, has increased dramatically over the last 5 years, and this trend is likely to continue. With this large infrastructure increase comes large installation costs and increased maintenance cost and complexity.

2.3.3 Networked Video and Audio

E-mail and the World Wide Web were the linchpins of the first two waves of Internet adoption. There is some reason to believe that a third wave will be inspired by the widespread availability of networked video and audio. While Internet telephony is currently a small part of the total telephony picture, increasing numbers of corporations are purchasing a new family of telecommunications switches incorporating seamless Internet telephony capabilities using existing analog handsets. These companies may provide a proving-ground for this technology: users accustomed to the virtues and vices of Internet audio transmission may be encouraged to adopt this technique in desktop and wireless computing.

Internet-based digital video has been long available, but has so far failed to have the expected impact. Existing technologies are complicated and generally deliver poor quality results. However, a combination of Moore's Law increases in computing power, increased network performance, and improved algorithms and software technologies appear to be starting to ameliorate these problems. Video is notoriously consumptive of network bandwidth: thus, widespread adoption of Internet video would require continued major adoption of new network infrastructure: gigabit LANs and 100Mb WANs are likely to become commonplace over the next 5 years.

2.3.4 Networking and the Digital Divide

One major barrier to the spread of networking in the academic community is the increasing Digital Divide [15] in the area of network services. The demand for "home broadband" (100-1000Kb/s bandwidth, ˜100ms latency) has increased dramatically, but deployment of broadband infrastructure has been slow. In Portland, some homes are served by Qwest's DSL service, and some by AT&T's Cable Modem service, but there are still many households with neither option. In addition, these technologies are still reasonably expensive both for installation and service, and the lead time and difficulty of installation is often large.

At the other end of the spectrum, "high-speed" dialup service is provided by a decreasing number of players: for example, the PSU Computer Science Department closed its dialin pool during the Summer of 2001. The remaining campus dialup services are generally time-limited, limiting their overall usefulness. The campus population thus increasingly drifts toward two other forms of home access: home broadband for those who have access to it and can afford it, and dialup Internet service through a commercial provider for the remaining users who can afford its somewhat lower (but still significant) monthly cost.

Such inhomogeneity is important: the academy is currently engaged in a large-scale effort, both at PSU and on the national level, to move education off-campus through synchronous and asynchronous distance and distributed learning. Technologies for synchronous learning generally require at least the latency and bandwidth of the typical home broadband connection to be feasible: even asynchronous schemes are greatly eased by sufficient bandwidth to obtain course materials and especially software.

2.4 Large-Scale Computing

With the growth of computing power discussed in Section 2.1, a whole new vista of possible applications in the academy naturally opens up. In addition to the obvious hard-science and engineering impacts, new disciplines and areas are beginning to be able to take advantage of powerful hardware and software. For example, biotechnology and the genetic sciences are increasingly dependent on computer simulation, modeling and analysis, and the cognitive sciences are being driven to some extent by an interaction between computational neural and learning technologies and traditional exploration of human cognition.

This raises a number of interesting questions. It is likely that some amount of parallel computing will always be a part of the academic mix: the constant factors afforded by parallelism allow early adopters to experiment with cutting edge applications “before their time”, and the new parallelism afforded by ubiquitous networking, cheap cluster nodes, and modern operating systems makes this sort of solution easier to implement.

In contrast to parallel computing, sequential computing is simple and cheap, and will continue to remain the dominant model over the next 5 years. However, one of the dominant models of computing from past generations, the very-high-speed sequential or vector computer, appears to be a casualty of the technology revolution: it is apparently no longer feasible to achieve computing speed increases without the ultra-large-scale integration afforded by microprocessor technology.

Advanced technologies, such as quantum computing and molecular or biological computing, would have huge impacts if successfully realized. However, they currently appear to be far more than 5 years from practical use.

The impact of these large-scale computing technologies is to some degree a matter of cost and availability. It is likely that some small number of parallel computing environments will continue to exist at PSU. These may be augmented by access to off-site resources, although problems with that model including availability, control, and reliability have up until now limited this option for most researchers. Fortunately, standard desktop machines are reaching performance levels where they can run all but the most demanding applications, and this will reduce the pressure for large-scale computing in the academic setting.

2.5 Environmental Issues

Two key events have brought environmental issues to the attention of the technology community recently. First, a report by a U.S. Department of Energy consultant [] claimed that electricity use by computing equipment is 7% of total national usage and expected to grow by 500% over the next ten years. While more detailed and careful analyses [] have exposed these figures as fundamentally flawed, the flawed figures still circulate, and are driving further efforts (beyond the long-standing EnergyStar [] efficiency program) to reduce computing power demands in the face of electricity shortages nationwide. In particular, the electricity problems of California’s Silicon Valley region [] combined with the huge concentration of computing equipment there has been was the source of much attention for a brief period in 2001.

Still and all, academic electrical requirements for computing equipment are unlikely to be particularly problematic in the short term. EnergyStar power-saving features work particularly well in the academic environment, where computer use tends to be highly intermittent.

A more important emerging problem is the recent recognition of the amount and toxicity of computer-related waste, a problem highlighted by the recent studies of lead contamination from discarded computers. In the academy as much as in industry, the need noted above to frequently discard almost all computing-related equipment produces a high-volume waste stream that must be dealt with. Currently, local and national computer recycling efforts are handling this stream effectively: it is likely that this trend will continue.

The largest potential impact of these environment threats is increased governmental and regulatory oversight of computing equipment deployment and disposal. This could place substantial new burdens on the administrative infrastructure to produce new reports and more effectively track hardware in the academic environment.

3 New User Interaction Technologies

Core computing is undergoing large-scale change. Concomitantly, user interaction technologies are evolving in fundamental ways. The extent of the user interaction revolution is only beginning to be felt: for most users, the majority of computer interactions still involve cathode-ray tubes, LCDs, and LEDs, keyboards, buttons, and knobs.

It would be difficult to even briefly touch on all of the user interaction technologies that are currently in the short-term pipeline. For now, a few of the most important examples must suffice: an output technology (Section 3.1), an input technology (Section 3.2), and a class of technologies from in-between (Section 3.3). Together, these three pieces may shortly begin to substantially transform the computer–user relationship.

3.1 The New Displays

Perhaps the most primitive and unsatisfactory feature of a modern computer is its display device. The cathode-ray tube (CRT) is a technology now more than 50 years old. The CRT is power-hungry (often consuming as much electricity as the entire rest of the computer combined), provides a poor-quality display (for example, resolutions of 150 dots per inch or better are considered absolutely minimal for text readability: reasonably-priced CRT-based monitors are not currently capable of this resolution), wastes large amounts of valuable desktop space, and is unreliable.

Nonetheless, the CRT is still the predominant mode of interaction with desktop computer systems, being inexpensive, well-understood and well-supported by existing hardware and software. It is also the case that innovation in this area has proven extraordinarily difficult.

Perhaps the only widespread alternative to the CRT in 2001 is the LCD display ubiquitous on laptops. While LCD displays are superior in many ways, they suffer from high cost, small viewable area, and fragility. Over the next five years, as these defects continue to be remedied, the LCD display should largely replace the CRT as the predominant desktop display device.

However, two even more important technologies are just now emerging, whose rapid development and deployment could entirely transform the display of digital information. Organic LEDs (OLEDs) [13] are carbon-based inks that can be printed onto flexible surfaces with an ordinary ink-jet printer. When electricity is applied to an OLED, it glows like an LED. The potential here is for high-resolution, low-cost, flexible displays that can be "tiled" to achieve large viewable areas.

Similarly, "Electronic Paper" [5] is a technology in which tiny machines attached to a flexible substrate change their displayed color in response to an electrical signal. Recent advances in nanotechnology have made it possible to fabricate such devices at a reasonable price. The potential here is for a paper replacement with cost, size, resolution, and display quality comparable to ordinary paper-and-ink, but allowing the contents of the page to change on the fly.

Both OLEDs and Electronic Paper are probably at least 3–5 years from widespread adoption. However, the success of either of these technologies would have profound consequences. Inexpensive roll-up computer displays would allow ubiquitous access to digital information, helping to close the digital divide, to alleviate the environmental impacts of massive use of paper and CRTs, and to change the whole nature of computer interaction. The consequences of being able to store, transport, and use a whole library in the form of a single electronic book printed on e-paper are almost unimaginably broad.

The probable schedule and potential impact of these new display technologies urgently needs careful assessment; while the technologies themselves may be deployed fairly rapidly, absorbing the impact of these technologies on the academy will surely be a lengthy and difficult process.

3.2 Language Understanding

For about 30 years now, researchers have been studying the problem of the recognition and understanding of language by computers [8]. Progress has been slow, for several reasons. Perhaps the most notable of these is that the term *recognition* is itself a proxy for a complex set of highly-interleaved tasks. The task of identifying, with moderate accuracy, particular phonemes in carefully inflected speech in a noise-free environment was solved long ago. But language is more than just speech, and an understanding of words, sentences, and ideas is crucial to accurate real-world speech recognition as well as to deeper language understanding.

Recently a technique known as Hidden Markov Modeling (HMM) in which recorded speech is compared to simulations of likely similar utterances has been used to make a great deal of headway in speech recognition in the absence of language understanding. It is now possible to dictate a letter to a computer and edit it only minimally to produce acceptable text.

There are several drawbacks to HMM technology. Perhaps the most important of these in the short term is the tremendous amount of computation required by this approach. The impact on the academy is substantial: the current generation of desktop hardware is barely powerful enough for this application, underscoring the need for a serious commitment to computer infrastructure if speech recognition is to become a regular feature of the educational environment.

In addition to speech recognition, other language understanding technologies are starting to play a role as well. In particular, gesture-based recognizers have been the subject of much investigation over the last few years. One inevitable longer-term consequence of the development of these technologies and the technologies of artificial intelligence is the integration of various input techniques to allow intelligent interaction with the computer.

3.3 "Intelligent" Interaction

Traditionally, artificial intelligence (AI) has been a somewhat stand-alone discipline. The quest for elusively-defined "intelligence" has been the purview of academic researchers whose applications have been in exotic areas such as robotics and computer chess.

In recent years, AI technologies [12] have broadened their power and scope. More powerful computers have made machine learning techniques such as neural networks and genetic algorithms feasible approaches to problem-solving. At the same time, traditional AI approaches such as heuristic methods for solving practical instances of hard problems have also benefited from increased hardware performance. Finally, an increasing appreciation of the complexities and subtleties of computer input and output as well as an increasing emphasis on real-time response and human-computer interaction have underscored the need for AI in day-to-day work.

The growth of AI in computer games is perhaps emblematic of this trend. Computer games of the 1990s relied primarily on high-quality 3D graphics and sound to win market share. As graphics hardware performance growth begins to decelerate, game programmers are beginning to focus on two critical but long-neglected areas: realistic world modeling and intelligent game play.

As a counterexample, Microsoft has experimented several times with *intelligent assistants* for their software products. Neither Bob nor Clippy was particularly useful or well-received, and there is no reason to believe that this sort of application will evolve into something usable any time soon.

What then, can we look forward to in intelligent interaction? Things like smart predictive typing to reduce the risk of repetitive stress injury. Speech recognizers and automated language translators of increasing quality are starting to become available. There is serious research into replacing the desktop metaphor with smart technologies for real-time interactive search and organization of information. Advances like these should have real impacts in the academic setting over the next five years.

3.4 The Growth Of Immersion

There was a time period in the mid-1990s when the media became obsessed with the notion of "virtual reality" (VR), the experience of immersion in a simulated world via goggles and gloves. It is fair to say that this sort of virtual-world immersion has not had much impact upon society in the past 5 years. There are several reasons for this. The obvious one is that the hardware and software needed to create an immersive virtual world is complex, expensive, and difficult to use and program.

However, expense and complexity are only two reasons why traditional VR has expanded slowly. Another important reason is that the usefulness and interest of VR lies largely in niche cases. For most experiences one might have, it is more realistic, easier, and cheaper to simply have them directly.

Another kind of immersion, on the other hand, has grown steadily. Interaction with virtual worlds using a standard desktop computer preserves many of the advantages of the more immersive technologies, while reducing the cost and complexity of interfaces. If it is possible to become immersed in the "virtual world" of a printed book, it is surely straightforward to become immersed in the much more interactive world of a computer game, as the popularity of massively-multiplayer online roleplaying games such as Everquest attests. Further, by providing more and more polished, interactive, and "realistic" interfaces, the graphical user interface of the modern computer has begun to blur the line between interacting with a VR world and interacting with a traditional software user interface.

It is likely that the use of sophisticated technologies such as real-time graphics and three-dimensional rendering hardware will continue to diffuse through the computing infrastructure. A modest investment in the hardware and software necessary to support this process may have large rewards in computing effectiveness.

4 The Changing Face Of Software

Predictions about software are even more difficult than predictions about growth of hardware technologies and infrastructure. Much of the current pattern of software development, deployment and use was completely unanticipated 5 years ago. Nonetheless, there are some trends that are both stable and important and thus deserve consideration even in this murky landscape.

4.1 Microsoft

Microsoft currently enjoys a monopoly [9] in most of the software categories used at academic institutions across the country. It is thus necessary to pay very close attention to the status and plans of the company over the next few years in order to react and plan appropriately.

Perhaps the biggest change at Microsoft in recent years is some small signs that sales are leveling off. Various competitors (notably open source, see Section 4.3) are once again making inroads in the traditional productivity suite applications: word processing, spreadsheet, and presentation. The operating system market is more complicated, but sales of the new Windows XP are in any case less than expected: this may be primarily due to fee increases and new licensing restrictions and technologies, which have discouraged many organizations from upgrading their existing Windows installations.

However, the likelihood is that Microsoft will remain the dominant player, if not the only player, in the operating system and productivity tools market over the next five years. There are several characteristics of Microsoft to consider when mapping a strategy here. First, Microsoft is litigious: it pays to arrange broad licensing schemes (even at extra cost) that preclude any accusations of illegal use. Second, Microsoft understands the value of broad academic penetration: this means that their software is usually available to academic institutions at reasonable cost and under reasonable terms. Finally, Microsoft tends to form one-sided partnerships: history suggests it would be wise to steer clear of any unusually close coupling between Microsoft and the academic institution, and to approach with extreme consciousness and awareness even normal partnership situations.

4.2 The Application Services Model

The history of software as intellectual property is an interesting topic, albeit too complex to discuss here. What is relevant is that, as of the late 1990s, a general consensus had been reached as to the terms of the agreement between software producers and consumers. The software producer sells (or gives away) the product under a license agreement known as an End-User License Agreement (EULA). The consumer then owns an unlimited-time right to use this version, but must pay for any desired upgrades to new versions. Since the software is licensed, rather than sold directly under copyright, the consumer typically gives up many of the rights associated with purchase of a copyrighted work, such as the right to resell it.

More recently, many software producers have been advocating a model in which consumers only license software in a time-limited fashion. In this Application Services Provider (ASP) model, the consumer pays for access to the latest versions of software packages on an as-needed basis. There are many variations on this theme, notably providing hosting service for the leased application and web-enabling the application.

The most prominent example of this is Microsoft's ".NET" initiative, which includes many components that are used on an ASP basis. In particular, the new Windows XP operating system has licensing features that some suggest signal a move toward an ASP model for the OS as well.

There are some potential advantages to the consumer in the ASP model. The consumer may be more easily able to find applications compatible with their platform and suitable for their niche-market needs. In addition, the process of accessing the software on-demand insures that the latest version of the software will be properly configured for the platform.

The downsides to the consumer are also very large. Not only does the amortized cost tend to be larger than in a licensed-software model, the consumer has much less cost control under the ASP model. The key issue here is that of "orphaned" data: if a user's documents are locked into a particular ASP application, then either an unacceptable change in licensing terms or unreliability in the ASP service can leave the user unable to access their own data. The obvious solution to this problem is to have multiple providers of ASP applications operating on common data formats: this is not possible for most types of data in the present environment, although XML (Section 4.4) may change that situation somewhat.

It currently appears unlikely that the ASP model will be widely adopted. However, this trend deserves close attention because of its possible large impacts on the academic software situation.

4.3 Open Source and Free Software

Folks have been writing software and giving it away for an extremely long time. Indeed, it was not until the early 1980s that it was entirely clear that binary-format computer programs are protected by U.S. Copyright [1]. Throughout the 1970s and 1980s, computer centers at academic institutions typically operated using a mixture of commercial software, "academic-use" software such as AT&T or Berkeley UNIX, and freely-available tools.

More recently, with the rise of the Free Software Foundation and the development of powerful free operating systems such as Linux and various versions of BSD UNIX, free software of various flavors has begun to be used significantly outside the core computing arena. The obvious advantage of Free Software is the savings in cost.

However, there are a number of secondary advantages that accrue to the use of free software in an academic environment: support costs tend to be low or zero, problems can be fixed directly rather than through external support, and the potential risk of expensive and difficult software "audits" and the like are minimized.

As a result, most computing outside the desktop arena is done either on a free operating system platform, or on a proprietary hardware-specific UNIX platform using free tools. The desktop, however, remains largely proprietary even in academic circles. The principal reason for this is the monopoly Microsoft (Section 4.1) has achieved in the desktop operating systems and applications market. It is unclear whether free software will achieve significant penetration in this area in the next 3–5 years; if it does so anywhere, however, it should be in the academic arena.

4.4 The Importance Of XML

The eXtended Markup Language XML [3] is a specification for a meta-language designed to address the problem of portable data in the modern era. Because XML is an ASCII-based format, documents are human-readable. Because XML is a structured format, computers can easily store and read complex relationships between data items. Because XML is a public standard, there is little potential for data lock-in. Over the last several years a number of types of data have begun to be routinely represented using XML.

Unfortunately, the rapid adoption of XML was spearheaded by the dot-com community: a community now in disarray and decline. As a result, several problems with XML have emerged as potential impediments to further adoption of this standard.

First, most of the structure of an XML document is governed not by the XML standard itself, but by a *Document Type Definition* (DTD) that describes the structure of a particular type of data. Without significant cross-vendor effort to standardize DTDs for a particular type of data, the interoperability advantages of XML are less achievable.

Second, the widespread adoption of XML is somewhat predicated on the easily availability of portable, powerful toolbases for manipulating and using it. The collapse of the dot-com world took with it many of the vendors of and customers for these sorts of tools.

Finally, XML is the victim of a classic bootstrap problem: widespread use can only be achieved if standard DTDs are developed and used by a wide range of applications: this will only happen if developers are convinced of the future widespread use of XML. Again, the collapse of the dot-com world made this appear somewhat less likely.

In spite of all of this, XML deserves to be widely adopted in academia: it solves a number of important problems in data representation, storage, manipulation, and interchange that are experienced daily by the academic community. Some attention should be paid to ways in which its use can be made more widespread in academic settings.

4.5 Intellectual Property "Management" Technologies

In the microcomputer world, in the late 1970s and early 1980s, the mantra of copy prevention technology was taken up in a major way by commercial software developers. The combination of weak legal protection for commercial software and convenient access to cheap duplication was felt to offer little disincentive to illegal use of software. An escalating arms race between the copiers and the copy-preventers led to increased software cost and decreased usability and reliability. In the end, a mixture of unprotected software and low-impact copy prevention schemes was adopted by most commercial software vendors: rather than observing a revenue decrease due to unauthorized copying, most vendors saw revenue increases due to the increased quality and decreased cost of their software.

The late 1990s and early 2000s appear to mark a similar era for intellectual property as a whole. For the first time, books, audio, video, and other forms of media are reproducible in perfect quality at little cost. In addition, software titles are becoming more expensive to produce, and software platforms are becoming more homogeneous. As a result, various sorts of copy prevention are starting to appear for all of these media. These schemes likely to continue to gain popularity for the next few years, although they are also likely to eventually tail off in popularity.

Table 1: Impact and Likelihood Of Changes

Change	**L**	**S**	**F**	**A**
Moore's Law (2.1)	H	M	M	M
Handhelds (2.2.1)	M	H	L	L
Home/Office (2.2.2)	H	M	H	M
Wireless (2.3.1)	M	H	H	M
Networking (2.3.2)	H	M	M	M
Networked A/V (2.3.3)	M	M	H	H
Digital Divide (2.3.4)	L	H	L	M
Large-Scale (2.4)	L	H	L	M
Environmental (2.5)	L	L	L	M
Displays (3.1)	M	H	H	H
Language (3.2)	M	H	H	H
Intelligence (3.3)	L	M	M	L
Immersion (3.4)	M	H	M	L
Microsoft (4.1)	H	L	L	M
ASPs (4.2)	M	M	M	M
Free SW (4.3)	H	M	M	M
XML (4.4)	M	L	L	M
Copy Prevention (4.5)	M	H	H	H

In the short term, there is a small potential impact to academia: the ability to have backup or archival copies of licensed media may be slightly more difficult. Although this is likely more of an issue for multimedia than for computers per se, specific classes of multimedia devices should be watched closely to avoid compatibility and usability problems.

5 Impacts and Conclusions

Table 1 summarizes the expected likelihood and impact of the changes described above. Each estimate is a qualitative value: either High (H), Medium (M), or Low (L). The table lists the lowest-level topics described in this document, together with the section in which each is discussed. For each topic, an estimate is given of the confidence (C) of the predictions of that chapter, and the impact of the predicted changes on each of three campus constituencies—Students (S), Faculty (F), and Administration and Staff (A)—if the predicted changes occur.

Like all such estimates, Table 1 must be taken with a large grain, if not an entire shaker, of salt. However, in the absence of better information, it at least gives some vague guidance about the importance of the changes predicted, and some minimal validation of the selection of topics in this report.

It is commonplace in the computing industry to draw analogies with the automobile industry. If these analogy hold, it would be reasonable to expect that the rapid growth and improvement of computing technologies will soon plateau, as the natural limits of the technology are reached.

In contrast, computing technology seems to be more like the discovery of fire, for which new implementations and uses are still appearing at a rapid pace thousands of years later. Much as "fire-based technologies" transformed human society completely, it is not unreasonable to expect computing technologies to continue as a transformative force on human society; a transformation to which the academy is hardly immune.

This report attempts to forecast some of the short-term changes to the academy that are both likely and transformative. As with all forecasts, it is likely to be increasingly inaccurate with time. Hopefully, enough accurate near-term detail is provided to assist academic decision-makers in the planning process.

Acknowledgements

Thanks to all those on PSU's Computing Staff and in PSU's Computer Science Department who had excellent comments and suggestions on earlier drafts. Thanks in particular to Jim Binkley for straightening out some of the authors' confusion about networking issues: many of the good ideas in the networking session are the result of this feedback.

References

[1] Marc Abrams. *World Wide Web: Beyond the Basics*. Prentice Hall, 1998. Chapter 9.

[2] Bluetooth SIG, Inc. Specification of the Bluetooth System, February 2001. Web document. URL `http://www.bluetooth.com/dev/specifications.asp` accessed January 21, 2002 02:45 UTC.

[3] Tim Bray, Jean Paoli, C.M. Sperberg-McQueen, and Eve Maler, editors. *Extensible Markup Language (XML) 1.0*. W3C, 2000. Online document. URL `http://www.w3.org/TR/2000/REC-xml-20001006` accessed January 21, 2002 02:45 UTC.

[4] Ariana Eunjung Cha. Dot-com collapse stifles innovators. *The Washington Post*, April 30 2001.

[5] Steve Ditlea. The electronic paper chase. *Scientific American*, 285(5), November 2001.

[6] Chris Gottbrath, Jeremy Bailin, Casey Meakin, Todd Thompson, and J.J. Charfman. The Effects of Moore's Law and Slacking on Large Computations, Dec 1999. Web document. URL `http://agave.as.arizona.edu/~chrisg/mooreslaw.html` accessed January 21, 2002 02:45 UTC.

[7] IEEE. IEEE Standard 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, 1997.

[8] Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics and Speech Recognition*. Prentice Hall, 2000.

[9] Stephen Labaton. U.S. vs. Microsoft: Appeals Court voids order for breaking up Microsoft, but finds it abused power. *New York Times*, June 29 2001.

[10] P. Messina, D. Culler, W. Pfeiffer, W. Martin, J. Oden, and G. Smith. Architecture—the high-performance computing continuum. *Communications of the ACM*, 41(11):36–44, November 1998.

[11] G.E. Moore. Cramming more components onto integrated circuits. *Electronics Magazine*, 38(8), Apr 1965.

[12] Nils Nilsson. *Artificial Intelligence: A New Synthesis*. Morgan Kaufmann, 1998.

[13] P.N. Powell. Scientists squeeze more light out of organic LEDs. *Laser Focus World*, 37(3), March 2001. Web document. URL `http://lfw.pennnet.com/Articles/print_screen.cfm?PUBLICATION_ID=12&ARTICLE_ID=95971` accessed January 21, 2002 02:45 UTC.

[14] R. S. Schaller. Moore's Law: past, present and future. *IEEE Spectrum*, 34(6):52–59, June 1997.

[15] National Telecommunications and Information Administration. Falling through the net: Defining the digital divide. Technical report, U.S. Department of Commerce, July 1999. Web document. URL `http://www.ntia.doc.gov/ntiahome/digitaldivide/index.html` accessed January 21, 2002 02:45 UTC.