The disruptive user — Internet appliances and the management of complexity

S E Gillett, W H Lehr, J T Wroclawski and D D Clark

Bringing networked computing to new users and new contexts entails a disruptive decrease in the level of user patience for complexity. This paper discusses the tensions involved in making devices as easy to use as traditional appliances, within the context of the open and rapidly changing Internet. It distinguishes class 1 appliances, whose function is fixed by the manufacturer, from class 2 appliance, whose functionality is determined by an associated service provider, and posits a third class of appliance that would achieve true ease of use by leaving control with the user while simultaneously automating much of the complexity associated with that control.

1. Introduction

Continuing growth in the communications market depends on an evolution beyond the desktop PC[1] in which networked computing capability — the Internet experience — is extended both to whole new classes of users and to new situations for existing users. Serving either of these markets, however, entails a disruptive decrease in the level of user patience for complexity. The former is the mass market — coveted as a large number of users, but demonstrating by virtue of late adoption that the current PC's balance of value delivered versus pain incurred is unacceptable. Also, while some existing users may be more tolerant of complexity in general, the trend is to have them use the Internet in environments in which they are likely to value simplicity more highly than before. These settings include non-work parts of the home (e.g. looking up a menu in the kitchen or playing on-line games in the den) and mobile environments such as the car. Even most existing users will only use the Internet in these new environments if it can meet their needs when they are in a hurry, in leisure space (whether mental or physical), or without the support of office staff.

The post-PC world offers us the chance to have computing pervade our lives, but users are unlikely to be enthusiastic unless the new applications and devices can finally deliver on the promise of enhancing and simplifying our lives. The disruptive user for the next phase of Internet growth is likely to be all of us who will resist the extension of computing beyond the PC and work/home office unless it is useful (i.e. we do not have to sacrifice key functionality), convenient and easy to use. This implies that future devices must appear to be much less complex than those that we have learned to accept in the PC world. With PCs, much of the flexibility is managed by user actions — users deciding what hardware peripherals to add, what software to load, or what configuration options to accept. This open environment is very friendly to the adoption of new services, but places strong requirements on the user to manage complexity. For the next generation of devices, much more of this complexity must be managed by the devices themselves, without direct user interaction. This is the sense in which we face the disruptive user.

So-called 'Internet appliances' have been touted as a way to address the needs both of power users for computing in new environments, and of new users put off by the complexity and limitations of the traditional PC platform. Quite a few new devices have sought to expand the markets for networked computing.

These include a number of devices like the Netpliance — essentially, scaled-down PCs that were supposed to allow Web cruising in the kitchen or family room via a simplified interface — and new mobile devices like the Palm VII PDA, Blackberry e-mail device, or cell phones with enhanced data capabilities (WAP phones, Nokia 9000-style phones, and so on).

Thus far, we see that the attempts to create 'computers for the rest of us' have failed miserably, while some of the devices targeted at power users have worked out. Among many reasons for this outcome are problems we identify with appliance strategies to date for achieving the true ease-of-use within the context of the Internet.

[1] In this paper 'PC' is used to mean personal computers in general, not just those running Windows.

2. Why is it so hard to make Internet appliances easy to use?

Why are plain old computing devices like PCs so hard to use? Because unlike most other devices we use in our lives — including objects like ovens, chairs, and cars — personal computers are designed to be general purpose. The hardware intentionally does not predetermine the computer's use — the software does. The software makes the computer flexibly upgradable and customisable. The user chooses what software to load and how to configure it after taking the computer home. It is true that computers increasingly come preloaded with the most commonly used applications. However, users still retain the ability to add new boards and peripherals, and load new software so that the PC can serve whole new purposes not necessarily envisioned by the PC designer. It is as if one could take a toaster, insert a floppy or CD, and get a microwave oven.

This general-purpose platform model dovetails nicely with the Internet, which was designed in the same way. Designers did not try to guess what users would do with it (i.e. what applications they would run); rather, they designed it to allow end users to choose how they wanted to communicate. In the process they enabled the dynamic we still see today in which new uses of the Internet emerge over time, e.g. the Web[2], audio and video streaming, peer-to-peer applications. In this environment, the PC was the perfect user device because it extended the flexibility inherent in the Internet to the end user.

Although this kind of flexibility has many desirable properties, including support for innovation [1], it incurs costs in complexity and unpredictability. New functions do not simply appear — users have to take purposeful action (i.e. download and install software) to enable them. The user's ability to decide which set of functions their PC will perform creates a combinatorially explosive set of configurations that cannot be systematically tested for interactions and performance, leading to the confusingly unpredictable behaviour we have all come to loathe about our PCs. Open systems give open possibilities, but these can quickly challenge the user with unwanted complexity. As with the old saw about VCRs that keep flashing '12:00' because their users cannot be bothered to wade through the manuals to learn how to set the time, most users never take advantage of the full flexibility of the PC because it is too complex to warrant the time to learn how to make it do what you really want.

Within this context, a trend towards 'Internet appliances' has emerged. The dictionary definition of an appliance is a device whose function is fixed to a particular purpose. Translating this notion into the world of software-defined functionality (the PC) and ever-new applications (the Internet) poses a daunting challenge. We consider to what extent restricting the functionality of the user device (making the PC more appliance-like) is likely to satisfy demand for ease-of-use.

One can think about Internet appliances in the same way as one might think about a real appliance — a device whose function is fixed at the time of manufacture. With this approach, which we call class 1, the manufacturer can resolve many (but not necessarily all) of the issues that introduce complexity for users in advance of the user's purchase of the device. Although some class 1 Internet appliances have been thought about, few have been built, for reasons that are discussed below (see section 3).

Alternatively, a class 2 concept describes an appliance with functionality that can change over time, but under the control of a service provider associated with the device, rather than the user. Here, the complexity is managed by a third party, transparently to the user. As we shall see, many class 2 appliances have been tried, largely without success.

From this discussion we will argue that a third class of appliances is needed, one which is much harder to achieve (it requires further research), but holds more promise for truly meeting the needs of users demanding simpler-appearing systems. Class 3 devices self-manage complexity by taking advantage of sophisticated adaptive and context-aware computing environments.

Although this one-dimensional classification of Internet appliances suffices to bring out several interesting issues which we discuss in the rest of this paper, it is worth noting here that reality can be even richer. As Fig 1 shows, the user's perception of the fixedness of a device's function is not necessarily the same as the physical reality of the device. Many different kinds of change can be made to appliances, and some of them — such as bug fixes automatically downloaded into a cell phone, or the translation of new e-mail formats (e.g. messages in HTML) into text — may never be noticed by users. The result is a grey area lying somewhere in between classes 1 and 2, and deserving of further consideration.

3. Class 1 — function fixed by manufacturer

Class 1 devices are the closest to traditional appliances like toasters and refrigerators. They can only do exactly what the manufacturer built them to do.

Portable scanners that can upload images over the Internet, cell phones enhanced with built-in Web browsers, and devices that play Internet radio stations are all examples of Internet appliances that could in theory be built as class 1

[2] Despite being almost emblematic of the Internet today, the Web was not an original Internet application. It is worth bearing in mind that without the properties of the general-purpose platform discussed in this paper, the Web could not have emerged to popularise the Internet roughly two decades after its invention.

		device is physically upgradable?	
		no	yes
user perceives change?	no	class 1: like traditional household appliances, e.g. toasters and washing machines	class 1b/2a: changes transparent to users, e.g. bug fixes, or translation of new formats into existing ones; or, class 3, if done well
	yes	class 3: if users notice lack of interoperability that emerges over time as Internet evolves	class 2: service provider adds new functionality, or changes existing function in non-transparent way

Fig 1 Internet appliance taxonomy.

devices (although most are not, in practice, for reasons we discuss below).

A class 1 Internet appliance has no capacity to change its function once it has been built. In contrast to the PC, it has no hardware or software customisability, no slots for cards that could modify or add to its functionality, and no capacity to accept new software. It may have system-level customisability if the manufacturer equips it with plugs to connect to other devices, but the nature of these plugs cannot be changed.

The class 1 approach can make appliances more user-friendly to the disruptive user in some ways. Because their hardware and software cannot be customised, class 1 appliances can be much more fully tested by the manufacturer. Furthermore, because a class 1 device's function is fixed for all time, the ways in which the device will be used can be better anticipated, simplifying the user interface design problem. Together, these factors could lead to appliances that are more reliable and intuitive to use than PCs.

On the other hand, users of class 1 Internet appliances are likely to suffer from lack of full interoperability with other Internet users, making the device's behaviour seem unpredictable and therefore frustrating to users, especially over time. This outcome arises from a fundamental tension between truly fixed functionality and the Internet's design for diversity and constant change. The Internet is diverse in a static sense — its 'hourglass' architecture provides for the co-existence of many different technologies and standards at any given time, both above and below the spanning layer[3].

Because of its flexibility, this same architecture facilitates the constant emergence of new functionality over time — both incremental additions to existing applications as well as entirely new applications.

By their design decisions, manufacturers control how interoperable their class 1 devices will be when they are introduced. Designers of e-mail-only appliances, for example, must trade off the cost and complexity of their devices against their ability to support different types of attachments. Should they support photos? How about files formatted in PDF or in Microsoft Word?

Will the resulting compromises be acceptable to users? — maybe, if they either never try to do anything beyond what the device supports or do not care if those functions do not work. But this static interoperability problem is only the beginning, because Internet change rapidly creates new applications and associated media formats with which to interoperate over time.

As we have argued elsewhere [1], the introduction of appliances is unlikely to slow this trend — or, in other words, rapid Internet change is exogenous. As a result, a class 1 Internet appliance is likely to become less interoperable and therefore more frustrating to use over time.

The potentially rapid obsolescence of class 1 Internet appliances suggests that they only work in the market-place if they are conceived of as devices that are frequently replaced[4]. One might reasonably question whether devices that require frequent replacement will really appeal to the techno-phobe or techno-indifferent user at whom class 1 devices would aim. Evidence from cell phones, however, suggests that this is not a serious stumbling block[5]. Consumers, it seems, are willing to replace fixed function

[3] In the architecture of the Internet, the spanning layer is exemplified by the Internet Protocol — the one layer that everyone absolutely must have in common. Layers above this include applications (e.g. e-mail, Web) and document formats (e.g. HTML, text, MP3, video formats such as Windows Media Player and RealNetworks), while layers below it include network infrastructure (e.g. Ethernet, ATM, wireless). Because these other layers, both above and below, can be 'wider' (i.e. support more than one alternative), this architecture is often drawn in an hourglass shape [2, 3].

[4] This observation has implications for environmental sustainability, suggesting that the use of recyclable materials is an important consideration in Internet appliance design.

[5] Data [4] shows that the average time to replace a cell phone in the USA went from once every four years in 1990 to once every two years in 1998. Worldwide replacement rates are estimated to be somewhat less frequent than in the USA, but still reasonably rapid (between 2½ and 3 years).

appliances reasonably quickly if there is a compelling reason to do so — and if the devices are relatively inexpensive.

This latter observation suggests a need for serious rethinking of Internet appliance product concepts that depend on expensive devices with fixed Internet functionality which will rapidly and visibly become obsolete. For example, a high-end refrigerator that tries to become a household nerve centre, by building in a screen and a fixed-for-all-time Web browser or e-mail client, is likely to have limited appeal[6]. Class 1 devices may have a place, but only if they are inexpensive, or their Internet functionality is sufficiently limited.

4. Class 2 — function fixed by service provider

Recognising the problems discussed in the previous section, most Internet appliance designers have taken a different approach, building devices that can be changed after they are built, typically by software download. In contrast to the PC model, with class 2 appliances the choice of what software to download, and when to install it, is not made by the user. A distinct feature of class 2 devices is their business model, which includes not only a device but also a service from a particular provider. The service may be basic connectivity (such as ISPs provide) but more often includes content-oriented features (such as America OnLine provides) as well.

Examples of class 2 Internet appliances abound — they include the Netpliance i-Opener (originally designed to work only with Netpliance's ISP-like service), AOL and MSN companion devices built by Gateway and Compaq (respectively), e-mail-only devices such as Landel Telecom's Mailbug and Cidco's Mailstation, personal video recorders such as TiVo's and SonicBlue's (formerly ReplayTV), and digital picture frames such as the Ceiva and Storybox. Each of these devices shares the characteristic that it only continues to work as long as the owner continues to subscribe to the associated service from the particular service provider. This service provider in turn determines the functionality of the box, by controlling what software runs on it and when, if ever, that software is updated.

Does the class 2 model make appliances easier to use? It does allow service providers to hide from users the complexity of adaptation to new functionality. The flip side, however, is that it shifts control over device functionality — in particular, the nature and timing of enhancements — away from users. A recent newspaper story recounted the surprise of a personal video recorder user who paused the picture to copy down a URL, only to have an advertisement pop up instead of the freeze-frame that would have been there the day before. Unbeknown to the user, the service provider had downloaded new software into the device and changed its functionality in ways that were desirable to the provider but not so warmly welcomed by the consumer.

While the class 2 model is better suited to rapid Internet change than class 1, users may still experience frustrations arising from incomplete interoperability. There are several reasons why providers may not always immediately update class 2 appliances with every possible innovation a user could want:

- low-cost devices may have technical limitations such as limited memory and disk space,
- service providers may decide that some innovations will add too much complexity to the user's experience, are irrelevant to the particular purpose of the appliance[7], or present legal or other risks (for example, enabling copyright infringement for which the provider may eventually be held liable).

Competitive strategy may also influence a service provider to delay the adoption of some innovations, especially if they involve interoperation with services offered by competitors. For example, the Compaq/Microsoft MSN companion appliances were introduced with support for multimedia files in Microsoft's Media Player format, but not in the (more prevalent) format promoted by their competitor, RealNetworks. Similarly, AOL-only appliances would probably not be quick to adopt interoperation with other services' instant message formats.

The line between technical limitation and anti-competitive abuse will be a fine one that, we speculate, consumers and policy makers will increasingly be called upon to distinguish with regard to Internet appliances.

The tepid market-place response to many class 2 appliances to date suggests that shifting control to service providers is not a panacea for ease of use. Although the class 2 approach offers the potential to hide complexity, it also happens to dovetail with marketing strategies that lock users in to particular services by imposing high switching costs. If class 2 appliances were much less expensive than PCs, consumers might not mind being so tied to a particular service provider. If they find they dislike the service at the start, or that over time it does not keep up with the Internet changes that they care about, or the associated service provider goes out of business, at least the user's sunk cost is

[6] See, for example, the Screenfridge [5]. Other concepts for Internet-connected refrigerators may prove more workable; for example, a fridge that uses the Internet only to report diagnostics about itself when a part fails is unlikely to cost significantly more or become obsolete much faster than it would without such functionality.

[7] For example, if new audio file formats emerge, digital picture frames will almost certainly not be updated to handle them. In contrast to the situation with PCs, there is no expectation that a picture frame appliance will be transformed into a stereo — just as there is no expectation that a toaster will be transformed into a microwave oven.

low. However, this is not yet the case, and it appears many potential appliance customers still choose the PC when faced with a cost-comparable choice between a complex, but general-purpose and more future-proof PC, and a simpler, but more limited and less adaptive class 2 Internet appliance.

The market experience with class 2 appliances to date suggests that they are more likely to succeed commercially if they find the happy medium that one industry contact described as '... the lightly walled garden with the well-marked exit' (Fig 2).

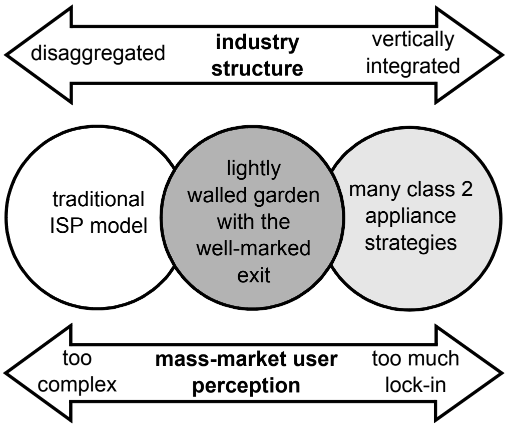

Fig 2 Spectrum of appliance business models.

This approach, exemplified by both America OnLine and DoCoMo, involves a level of service integration that is enough to get the user started — or to provide a complete enough experience for non-adventurous users — but does not close the door to additional content. DoCoMo, for example, selects a relatively small number of content sites[8] to be 'official', meaning that they receive preferential navigational treatment (e.g. menu placement) and can integrate their billing with DoCoMo — but they also provide a platform through which users can reach a much larger number of unaffiliated sites.

5. Class 3 — the Holy Grail

Achieving true ease-of-use requires attention to the particulars of what makes PCs and the Internet frustrating to use. Good solutions to real user problems may not exist today, but if they did, we could posit a third class of appliances.

Class 3 appliances would be like class 2 in that devices could be updated, but unlike class 2 and like PCs in that users would control the updates. Class 3 appliances would differ from PCs, however, in requiring less sophistication and involvement from the user to accomplish updates. They may also support a more limited range of functions than a PC, much as a Palm Pilot currently functions as a sort of scaled down PC[9].

Rather than requiring users to take explicit actions to keep their devices current with every possible upgrade, a class 3 appliance might be equipped with software that is intelligent enough to observe and interpret user behaviour, determine when an upgrade or enhancement is needed, and make it happen. PCs are already taking primitive steps in this direction, such as automated installation of operating system updates and dialog boxes that automatically appear, asking users if they would like to download the plug-in software they need to view content in a newly released Web format.

The development of methods to more broadly, reliably, and proactively simplify the user experience, through this kind of implicit reconfiguration of the appliance, is an active area of academic and industry computer science research. The focus is on leaving control with the user while simultaneously removing much of the complexity associated with that control through automation. If this research is to succeed in encouraging user acceptance of complex devices, it must also consider issues of security, device integrity and user privacy as well as the more obvious technical challenges.

6. Conclusions

To date, most Internet appliances have tried to achieve ease by closing up open systems. This strategy appears to be based on the premise that reducing choice, hiding heterogeneity and constraining functionality will make Internet appliances easier to use than PCs, for new and power users alike.

This strategy has achieved little market-place success. The Internet is an open system and here to stay. At least with the current state of the appliance market, appliances need to adapt to it, rather than the other way around. The interoperability challenges raised by the class 1 and class 2 models discussed above are fundamental.

The real challenge, then, is not to close the system but to find ways to make an open system easier on its users. Research is under way in many laboratories, including BTexact Technologies at Adastral Park, as well as MIT's Laboratory for Computer Science, on the specifics of how to do just that. In particular, this will involve observation of user behaviour to discern between those decisions that are and those that are not important to any individual user, making it possible to reduce the number of active decisions each user has to make while maintaining the choices that they demonstrate, over time, they care about.

[8] About 1100 in February 2001, versus 24 000 unofficial sites [6].

[9] Although, in principle, a class 3 device could have at least as much functionality as a class 1 or 2 device, it is possible that certain elements that may be critical to the deployment of specific classes of applications may appear first in class 1 or 2 devices (e.g. GPS for location-dependent services).

Acknowledgements

The author gratefully acknowledges the support of the sponsors of the MIT Program on Internet and Telecommunications Convergence (ITC), including BT; details are listed on ITC's Web site [7].

References

1 Gillett S E, Lehr W H, Wroclawski J T and Clark D D: 'Do appliances threaten Internet innovation?', IEEE Communications, 39 No 10 (October 2001).

2 Lehr W and Kavassalis P: 'The flexible specialization path of the Internet', in Bohlin E, Brodin K, Lundren A and Thorngren B (Eds): 'Convergence in communications and beyond', Elsevier Science, Amsterdam (2000).

3 Computer Science and Telecommunications Board, National Research Council: 'Realizing the Information Future: The Internet and Beyond', National Academy Press, Washington, DC (1994).

4 Cahners Web site — http://www.cahners.com

5 Screenfridge — http://www.electrolux.com/screenfridge/

6 MobileInfo — http://www.mobileinfo.com/

7 MIT ITC program — htt://itc.mit.edu/

Sharon Gillett is a Research Associate at the Massachusetts Institute of Technology (MIT) and Executive Direct of the MIT Program on Internet and Telecoms Convergence. She is also the Principal of Victory Research, a consultancy. Her research focuses on Internet infrastructure technology and policy. Her particular areas of expertise include broadband access to the Internet, the post-PC Internet, and universal service. She received her MBA and MS in Technology and Policy from MIT in 1995 and her AB in Physics from Harvard-Radcliffe in 1982. From 1982-1992 she developed computer networking software and managed projects for BBN Communications and Thinking Machines Corporation.

William Lehr is a Research Associate at the Massachusetts Institute of Technology and Associate Director of the MIT Program on Internet and Telecoms Convergence (MIT ITC). He is also an associate research scholar on the faculty of Columbia University's Graduate School of Business where he was a professor from 1991 to 1996, and a research associate at the Columbia Institute of Tele-Information. He teaches courses on the economics, business strategy, and public policy issues facing telecommunications, Internet, and eCommerce companies. In addition to his academic research, he provides litigation, economic, and business strategy consulting services for firms in the information technology industries. He holds a PhD in Economics from Stanford (1992), an MBA from the Wharton Graduate School (1985), and MSE (1984), BS (1979) and BA (1979) degrees from the University of Pennsylvania.

John Wroclawski is a Research Scientist with the Advanced Network Architecture Group at MIT's Laboratory for Computer Science. His technical interests include core principles of self-organising systems, the architecture, technology and protocols of large, de-centralised communication systems such as the Internet, and systems aspects of pervasive computing. His current work includes the Personal Router, aimed at altering the economics of ubiquitous wireless access, and the NewArch project, focusing on Internet architecture for the 10 to 20 year timeframe. He has contributed to several generations of research in the field of network quality of service, and has authored and co-authored a number of standards for IP QoS support. He is a member of the IEEE, ACM and the IETF, where he chairs the ISSLL working group and is a member of the Transport Area directorate. He holds SB and SM degrees from MIT.

David Clark is a Senior Research Scientist at the MIT Laboratory for Computer Science, where he has worked since 1973, when he received his PhD there. Since the mid 1970s, he has been leading the development of the Internet; from 1981-1989 he acted as Chief Protocol Architect in this development, and chaired the Internet Activities Board. Recent activities include extensions to the Internet to support real-time traffic, explicit allocation of service, pricing and related economic issues, and policy issues surrounding the Internet, such as local loop deployment. His current research looks at new models for deployment of wireless communications in the post-PC device era, and re-thinking of the architectural underpinnings of the Internet. He has also worked on computer and communications security. He is the Principal Investigator of the MIT Program on Internet and Telecoms Convergence. He is also a member of the National Academy of Engineering and chairman of the Computer Science and Telecommunications Board of the National Research Council, where he has contributed to a number of studies on the societal and policy impact of computer communications. He is a Fellow of the ACM and the IEEE.