Smart and stupid networks: Why the Internet is like Microsoft

Andrew Odlyzko
AT&T Labs - Research
amo@research.att.com

Revised version, October 6, 1998.

Abstract

Is the Internet growing primarily because it is a dumb network, one that simply delivers packets from one point to another? Probably not. If were a dumb network, we surely would not need huge and rapidly growing ranks of network professionals. A more detailed look suggests that the Internet is succeeding largely for the same reasons that led the PC to dominate the mainframe, and are responsible for the success of Microsoft. Like the PC, the Internet offers an irresistible bargain to a crucial constituency, namely developers, while managing to conceal the burden it places on users.

1. Introduction

The Internet is growing explosively, and is even threatening to take over transport of voice calls. Popular press often explains the rise of the Internet as a result of its greater efficiency in using transmission facilities [Keller]. Another explanation, popular among computing and communication experts, is that the Internet reflects the migration of intelligence to the edges, and thus leads to a dumb network that just transports bits (cf. [Isenberg]).

While both of these explanations are appealing and have evidence supporting them, they are not entirely satisfactory. Consider costs. While packet data transmission is likely to eventually become much less expensive than the circuit switched network, today most corporations spend more to send large files over their packet networks than they would if they used modems on the public voice phone network [Odlyzko]. As for the intelligence of the network, note that while the processing power of computers connected to the Internet is growing rapidly, there has always been more intelligence at the edges of the network than inside it, in the human brains at the ends of a connection. Thus the popular explanations are at best incomplete.

My thesis is that many of the factors powering the ascent of the Internet are similar to those that led the PC to displace the mainframe and were exploited by Microsoft to dominate PC software. In

both networks and computing, the end users typically care just about getting a few crucial tasks done. However, end users are not necessarily the most important players. In rapidly changing fields, application developers are the crucial ones, creating tools that attract users, tools that users do not know ahead of time they need. Both the Internet and the PC offered developers superior platforms with low intellectual and financial barriers to entry. The Internet won because it could be treated by developers as a dumb network, one that simply moves packets around. This led to an outpouring of creativity, with individuals or small groups creating "killer apps," as Andreessen and Bina did with Mosaic. In PCs, a similar phenomenon applied, with Bricklin and Frankston creating VisiCalc and later Kapor and his group creating Lotus 1-2-3. Microsoft achieved its dominance because it was better at catering to the developer community than the competition.

The Internet and the PC (especially Microsoft with its operating systems) were superb at serving developers, and reasonably good at serving the early adopters, the "power users." However, they were not optimized for the bulk of end users. Little attention was paid to human factors. The result is that both networking and computing are frustrating for end users. The complexities of competing systems (networks like CompuServe and Microsoft Network, and the mainframe in the computing arena) were lessened, but were not eliminated. Instead, most of those complexities were tossed into the laps of computing and networking support staffs. Even then, the end users could not be spared completely. We can see the results in the statistics of household penetration. In spite of all the hoopla about the Internet, only about half of U.S. households that have PCs (and thus about a quarter of all households) also have Internet connections, a much smaller fraction than have phone or cable TV service. Since Internet accounts are available to most of the U.S. population for about $20 per month, cost is unlikely to be a major deterrent, especially since most PC owners have spent upwards of two thousand dollars on their machines. Clearly there is something lacking to the Internet's appeal that so many millions of PC users stay away from it.

The rise of the PC and the Internet has resulted in a migration of intelligence towards the edges, in both networking and computing. Unfortunately this migration has also led to the migration of administration and maintenance duties towards the edges, and this has forced wasteful duplication of effort. I suggest that much can be learned from the experience with the mainframe and the phone network that would help alleviate this problem. I do not advocate the "Intelligent Network," with all functionality provided from inside. However, we could gain if some of the intelligence that is now required to make the Internet function either were pulled inside the network, or else resided at devices at the edges of the

network that were administered centrally. Intelligence would still be primarily at the edges, but much of it would be invisible to users and would lessen their support burdens.

The ideal interface should look simple to users, and conceal the intelligence within. By that standard, neither the PC nor the Internet is dumb enough. What systems do people get enthusiastic about? The Palm Pilot is an excellent example. It has sophisticated software and hardware, but that sophistication is invisible, and makes the few important tasks go smoothly. That level of user friendliness is what we should be aiming for. I advocate making the Internet appear even simpler than it does now, more of a "stupid network," by putting more functionality inside the network, and foregoing many of the Quality of Service developments that would require increased work for system administrators. This might lead to lower utilization of the physical facilities, but should lower total costs.

Finally, a disclaimer. The analogy between the Internet and Microsoft is not meant to be carried too far. For example, the arguments here say nothing about Microsoft's business practices, nor about the merits of the antitrust case against Microsoft. Only some technological, economic, and sociological factors that affect the spread of new technologies are conside

2. What do people want?

People want simplicity and stability, but they also want flexibility. The IBM mainframe, just as the voice phone system, stressed stability. The current era has other priorities. In an insightful piece in the *New York Times*, Edward Tenner [Tenner] points out that

> Microsoft has triumphed because it has given us what we asked for: constant novelty coupled with acceptable stability, rather than the other way around. Microsoft encouraged our impulses to embrace fashion, affirm conformity, love planned obsolescence. People talk simplicity but buy features and pay the consequences. Complex features multiply hidden costs and erode both efficiency and simplicity.

Although there are widespread warnings not to buy any Microsoft product until at least the third upgrade, they are widely disregarded. There are even people who pay to be beta testers, helping producers debug the producers' software. Clearly there is a large constituency for the kinds of products Microsoft and other software companies produce, buggy and user-unfriendly as they are.

In communications as in computing, people want networks that are both smart and stupid. Ideally the network should be stupid in the sense of having a simple interface, yet be smart enough to do

what the user wishes. The original telephone system had these qualities. Unlike the current automated version, the original had operators handling all calls, operators who could be asked to do a variety of chores. The functionality provided by those operators was valued enough to generate resistance to direct dialing.

Vestiges of the original phone system are still with us. Those lucky enough to have access to the White House switchboard rave about the service they get. Armed with finely honed skills, a large database, and the magic words "This is the White House calling," the professionals there are reputed to be able to find anyone, anywhere, at any time. This is the ideal stupid/smart network, with a simple interface that understands even spoken commands and interprets their nuances.

The transition from manual operations to the direct dial voice phone network led to a loss of the flexibility and expertise that human intelligence inside the network provided. This transition did lower the costs, though, and it did provide simplicity. The phone network is the most ubiquitous communication system so far, and it does allow anyone in Tulsa to call Timbuktu with a minimum of fuss. It has been adopted for uses far beyond the early predictions [deSolaP], and it has become a crucial part of all industrialized societies, much larger in revenues than even the airlines, and far larger than the entertainment industry. However, the traditional phone's lack of flexibility was an annoying constraint, especially as technology provided astounding progress in other areas, especially computing. The phone network required sophisticated technology to design and operate, but this technology was esoteric and invisible and did not allow users much control.

The deficiencies of the traditional phone system were partially remedied by putting more intelligence into the core of the network (with features such as call waiting) and by attaching intelligent devices at the edges (such as answering machines). These attempts were successful in providing more flexibility, but at huge cost both in development of the new features, and in simplicity of operation. Nobody enjoys getting trapped in "voice-mail purgatory," or having to keep track of a dozen or more numbers (phone, fax, pager, access code, phone credit card, etc.), many of which change every few months, or remembering all the commands ("Is it *68 or *86?") needed to make the system work.

Compared to the voice phone network, the Internet is a far more flexible medium. It has been described as "cheap and stupid because it was being built for smart endpoints called computers" [Petzinger]. Unfortunately it is neither cheap nor stupid. If it were, would we need hordes of network experts fluent in the language of BGP, caching, DNS, bandwidth management, Layer 3 switching, dual homing, firewalls, proxies, and other arcane topics? Would we also have the flourishing communica-

tions outsourcing business that AT&T Solutions, Andersen Consulting, and other companies enjoy?

3. What can people be made to want?

It is often said that “the people” or “the market” demand a particular product or feature. This may be appropriate for established categories such as cars, but less so for rapidly evolving fields such as information technologies. How much demand was there for a Web browser before Andreessen and Bina created Mosaic? The secret of success in the communications and computing areas has been to divine what people might want to use. In that environment, the dominant role belongs to those who create new products.

How can innovators create markets for their products and services? The easiest way to introduce a new technology is as a substitute for another one, a substitute that is better in cost or features. The jet engine replaced the turboprop that way. However, such simple substitutions are possible only in some markets. In general, it is hard to dislodge an established technology. It is far easier to develop unique new applications. The PC did not gain primacy over the mainframe through a frontal assault, by taking over banks' data processing centers, say. Neither did the Internet overshadow the switched voice network by carrying voice calls. Mature technologies usually are well adopted to their main task. It took over almost a decade from the introduction of the IBM PC, which set the industry standard, until it became accepted wisdom that the PC would be dominant. Even now, after a reinvigorating shot of PC technology, the mainframe continues to fill a substantial market niche. The most ambitious goals that the IBM executives had for their PC creation seem laughably modest today. However, those goals were set in a world dominated by mainframes, where the role of a PC was indeed limited. It took the development of new applications, especially spreadsheets like Lotus 1-2-3, to power the growth of the PC market.

Similarly, the Internet was dominated initially by email, and more recently by the Web, services that the traditional phone network could not provide. Furthermore, while the Internet has grown at astounding rates, it has done so largely by utilizing the infrastructure of the phone network. Twelve years ago, the NSFNet backbone (which evolved into the current Internet) consisted of a few dozen 56 Kbps circuits, of the kind that can carry a single phone conversation. Even today, the Internet is still considerably smaller than that phone network. Despite its fantastically rapid growth rate, the Internet carries little voice traffic, and its growth has come from the development of novel services.

As the Internet matures, it is worthwhile to see whether it can benefit from the lessons of the

development of the phone system. After all, the phone also started out as a niche product, unworthy of the attention of the Western Union telegraph giant. Initially it was an extremely expensive service for limited uses. In 1896, the basic monthly fee for a phone connection in New York City was $20. A century later, an Internet account is also $20 per month, but the purchasing power of this $20 is vastly different. In 1896, $20 was more than half a month's pay for a worker, comparable to about $1,000 today. How many Internet users would there be at $1,000 per month? To attain its ubiquity, with over 700 million lines around the world, the phone network had to lower costs and become sufficiently user friendly for the vast majority of the population to accept it.

4. Platforms, users, and application developers

When the search is on for compelling new products, it is the potential creators of these "killer apps" who have to be catered to. From the perspective of the innovators, the shortcomings of the mainframe and the phone network were that applications were extremely hard to create, requiring extensive knowledge of these platforms and expensive equipment, neither of which could be possessed by the proverbial teenager in a garage. Modern mainframe and phone switch development teams employ hundreds of software experts for many months at a time. The PC offered a better platform for applications developers than did the mainframe, and Microsoft offered a better platform than did competitors such as Apple. Small groups or even individuals could create novel applications with minor investment. This is crucial when the success of an application cannot be predicted beforehand, and most efforts fail.

It is worth noting that Microsoft's revenues have stayed around 10% of those of the entire PC software industry. It is true that Microsoft has taken the lion's share (around 30%) of the profits in that industry, and that other companies tend to oscillate between hope of being bought by the Redmond giant and fear of being crushed by it. However, that is not directly relevant to this essay. The point is that Microsoft succeeded by catering to developers, and its success has rested on their ability to craft new products and services on the Microsoft platform. Microsoft's victory over Apple, and later over IBM with its OS/2 operating system, owed much to gaining the support of developers.

The success of the Internet was also due largely to its offering a superior platform to application developers. Remember the introduction of the Microsoft Network, and the fear that it would dominate communications? It simply could not compete with the Internet, with its open standards and economies of scale. The big attraction of the Internet to developers was that it could indeed be treated like a dumb network, one that just carried packets from one point to another. Often only minimal knowledge of

TCP/IP was required to develop products for the Internet.

The PC and the Internet have benefited from the clean functional differentiation among service layers. This has allowed specialized players (Microsoft and Intel being the obvious examples) to exploit economies of scale and network effects to greatest advantage. However, this specialization has created its own costs, which we will consider in greater detail below. As a result, it is not clear if there were any real costs savings in carrying out traditional tasks. The main advantage of the Internet and the PC appears to have been in stimulating creativity among applications developers.

Of course, new products do have to be accepted in the marketplace. However, it is not the reaction of the vast majority of potential users that matters the most in the early stages of the development of a new product or service. For application developers, the first task is to hook the "early adopters," the people who are sophisticated and willing to devote time and energy to learning new tools and fitting them into their environment. Once these folks are on board, one can then worry about the mass market, where human factors are more important.

5. The hidden and not so hidden costs

The PC and the Internet are great creations that have stimulated creativity, invigorated the economy, and represent the future. However, they carry substantial hidden burdens. As an example, Microsoft gives developers the freedom to overwrite DLLs at will. The hapless users who suddenly find that crucial programs no longer work have to figure out what went wrong, in what order to reinstall all the packages, etc. Usually they end up searching for somebody knowledgeable to do it.

A prominent industry analyst likes Windows98 because he thinks it will cut in half the 26 minutes per week he spends booting up and shutting down his Windows95 PC [Dodge]. If other owners of the more than 200 million PCs save similar amounts, then, at a modest average valuation of their time of $15 per hour, the annual savings will come to over $33 billion, more than twice Microsoft's annual revenues! Yet, so far, such factors have not been a serious concern for Microsoft.

The main hidden costs of both PCs and the Internet are in the support operations. A Gartner Group study estimates that in a PC/LAN environment, total cost of ownership is 80% labor and 20% capital, while in a data center, the percentages are reversed [KirwinC]. As a result, both PCs and the Internet are not "cheap and stupid" [Petzinger], but expensive and intelligent, because they all require plenty of time from smart human brains to make them work effectively. These costs have been bearable in the

past largely because they were not visible, often carried by some support organization.

Rapid change always carries substantial costs. Schumpeter, the apostle of "creative destruction," championed capitalism not for efficiency (as, lacking our extra half a century of experience, he thought socialism could be more productive [Schumpeter]), but for its innovation. The problem is how to lower the costs that innovation carries.

6. Conclusions, recommendations, and predictions

The computing and networking industries are changing, and are paying more attention to their hidden costs. One reason is that those costs are getting harder to hide, and corporations are waking up to the $10,000 that it costs a company to operate a PC for a year. That is why the idea of a network computer (NC) was received so enthusiastically. Even if the NC does not make substantial inroads, it does appear to have frightened Microsoft and other software producers into paying more attention to total system costs. They are working to simplify administration of their systems. Typically this involves profligate use of resources, such as hundreds of megabytes of disk space for software installations. As the saying goes, "What Andy Grove giveth, Bill Gates taketh away," with the hardware provided by Intel quickly filled by bloated software from Microsoft. From a total system point of view, that is the preferred way to go. Economies such as using two digits to designate years are more likely to harm than help, when resources grow at an exponential pace.

Another reason to expect change is that growth in computer users is slowing down. To increase household penetration of PCs and the Internet, it appears that greater simplicity is desirable (and simpler devices such as WebTV may be one way to achieve it). This will require intelligence primarily at the edges of the network, not inside it, but intelligence that does not require extensive involvement by the users or systems experts to make it work.

Although there are forces that are pushing the IT industry towards simplification, other forces are working in the other direction. The industry is still searching for the next "killer app," and application developers are in control. Further, attempts to provide differentiated service levels on the Internet are complicating the scene for both developers and users. Whether the result will be more or less complexity is hard to predict. However, the fundamental attraction of a dumb-looking network and dumb-looking computers is undeniable, and the most successful companies are likely to be the ones that can deliver in this area. The trouble with the PC and the Internet is that they are not dumb enough. We should gain from putting more resources and intelligence into computing and networking to make

them seem dumber.

Acknowledgements: I thank Phil Agre, Dave Belanger, Marjory Blumentha, Art Deacon, Bob Frankston, Amy Friedlander, David Isenberg, Jeff Lagarias, Ed Tenner, Ed Vielmetti, and Bill Woodcock for stimulating correspondence and conversations.

References

[deSolaP] I. de Sola Pool, *Forecasting the Telephone: A Retrospective Technology Assessment,* Ablex, 1983.

[Dodge] J. Dodge, A contrarian's view of Windows 98: I like it, *PC Week,* April 27, 1998, p. 3. Available at ⟨http://www.zdnet.com/pcweek/news/0427/27week.html⟩.

[Isenberg] D. Isenberg, Rise of the stupid network: Why the Intelligent Network was once a good idea, but isn't anymore. One telephone company nerd's odd perspective on the changing value proposition, *Computer Telephony,* August 1997, pp. 16, 18, 20, 24, 26. Available at ⟨http://www.manymedia.com/david/stupid.html⟩. Revised version, The dawn of the stupid network, in *ACM netWorker,* vol. 2, no. 1 (Feb.-Mar. 1998), pp. 24-31.

[Keller] J. J. Keller, Ex-MFS managers plan to build global network based on Internet, *Wall Street J.,* Jan. 20, 1998.

[KirwinC] B. Kirwin and D. Cappuccio, Technology migrations: Toting up the hidden costs, *Data Communications,* May 1998, pp. 31-32. Available at ⟨http://www.data.com/business_case/migrations.html⟩.

[Odlyzko] A. M. Odlyzko, The economics of the Internet: Utility, utilization, pricing, and Quality of Service. Available at ⟨http://www.research.att.com/∼amo⟩.

[Petzinger] T. Petzinger, Jr., A smart business model is seen in a "Stupid Network," *Wall Street J.,* Feb. 20, 1998, p. A1.

[Schumpeter] J. Schumpeter, *Capitalism, Socialism, and Democracy,* Harper, 1942.

[Tenner] E. Tenner, The Microsoft we deserve, *New York Times,* Op-Ed page, March 8, 1998.