

# Nathan - Roadkill

## InterOffice Memo

**To:** List

**From:** Nathan P. Myhrvold

**Date:** September 8, 1993

**Subject:** **Road Kill on the Information Highway**

---

Technological changes often have enormous consequences. Microsoft has been the beneficiary of this effect for the last 17 years, and a variety of other companies in the computer industry have either had their day in the sun, or have fallen past the crest of the wave and have suffered as a result. Although the shifting fortunes of companies within the computer industry are naturally quite important to those of us who are participants, the effect on the world at large has actually been rather modest. The confluence of wide area digital communications and ever cheaper computing is going to be a *lot* more traumatic and far ranging than PCs have been. This memo is about some of those changes and how they will effect a number of industries.

Many of the ideas described here spring from numerous conversations with people at Microsoft and elsewhere. I'd thank everybody individually, but alas since my life isn't yet on line its hard to enumerate them all. I would like to thank Cindy Wilson for confirming a number of facts for me and providing terrific library support.

## A Computer On Every Desk

Personal computers have become incredibly useful tools - whether for individuals or organizations. The world *writes* with PCs - whether composing a simple letter, writing a novel, laying out *Time* magazine or drafting the blueprints for new building. We make *decisions* with personal computers - analyzing an investment or creating a budget and asking "what if?". These are all very valuable tasks, but it is also important to keep some perspective on the scope of the changes. Word processing has replaced the typewriter - but outside of the typewriter business one could ask "so what?". Spreadsheets, one of the strongest applications categories have replaced the *columnar pad* and the adding machine. The spreadsheet implementation of a columnar pad is a lot more convenient and has broadened the user base to some people who weren't big columnar pad users,

but this is the basic functionality. I remember seeing those funny pads with all the columns in the back of stationary stores years ago. I recently checked, and found that they, along with all those delightfully bizarre forms of K+E graph paper, are still carried in a good stationary store (albeit at reduced volumes). Despite the enormity of the personal computer "revolution" within the parochial confines of the computer industry, it really hasn't been that much of a revolution in society as a whole. How radical can a revolution be if its rallying cry (at least implicitly) is "Death To Columnar Pads!" ?

Another way to see this is to look at our own mission statement - *A computer in every home and on every desk*. This was a highly unconventional vision in the context of water cooled mainframes humming in the machine room, but it's really rather modest in the larger context of society as a whole. In actual practice, the mission that we and the rest of the industry have actually delivered is somewhat shorter than way we normally phrase it, it's really - *A computer on every desk*. We have been pretty good about getting computers onto desks whether they are in the office, or in the den at home, but desks are still our primary strength. We've done a bit better in some areas than others. Laptops (like the Omnibook I'm using right now on United 871 to JFK) have allowed people to take their desk activities with them. This is balanced by the fact that there are many real desks, such as those in the school classroom, that haven't been populated with computers.

The metaphor of the desk top which inspired researchers at Xerox PARC has largely been realized - personal computers and the software they run are primarily dedicated to service the activities that people do at a desk. Most of the product development done at Microsoft is focused on making incremental improvements to this basic mission, or in integrating the last remaining desk top activities. Our Microsoft AtWork campaign will unify PCs with the telephone, FAX and copier, closing on the last big gaps. However challenging (and profitable) this may be, we must not forget that our industry has grown up in a rather restricted environment. It is if we lived in a world where the only furniture was a desk, and columnar pads were on every bestseller list.

## The Importance of Being Exponential

I was recently interviewed by a guy doing an article to commemorate the 40th anniversary of *Playboy* magazine. He wanted to know what computing would be like 40 years hence, in the year 2033. This kind of extrapolation is clearly fraught with difficulty, but that doesn't mean that we can't try our best. In the last 20 years the overall improvement in the price/performance ratio of computing has been about a factor of one million. There is every reason to believe that this will continue for the next 20 years; in fact, the technological road map appears reasonably clear. This will yield another factor of one million by 2013. It is hard to predict what technology will be promising then, but I'm optimistic enough to believe that the trend will continue.

Laboratories are already operating "ballistic" transistors which have switching times on the order of a femtosecond. That is  $10^{-15}$  of a second, or about 10 million times faster than the transistors in this year's fastest microprocessor. The trick is simple enough in principle - reduce the size of the semiconductor component and the current flow so that the electrons don't bump into each other or the semiconductor atoms. In addition to being much faster, this dramatically reduces power drain and heat dissipation. The next stage is more of the same - people are currently experimenting with the "single electron transistor", in which a single bit is represented by a lone electron. This is not only very fast, it is also the ultimate in low power computing!

This raises a couple of other amusing points. Every advance in speed makes computers physically

smaller. Switching speed in semiconductors is directly related to their size, and at another scale, the delay or latency caused by time for a signal to travel from one part of the computer to the next is limited by transit time at the speed of light. At the speed contemplated above this is extremely significant. If you have a computer with a femtosecond cycle time, then it takes about 1 million CPU cycles for a signal to travel one foot. As a point of comparison, a hot processor of 1993 with a 100 MHz clock rate (10 nanosecond cycle time) would have a similar relative wait time in terms of clock cycles if it was sending a signal about 1860 miles. The latency associated with going across the country today will occur in moving across your desk in the future. The amazingly fast computers of 2033 will of necessity be *amazingly small* because you can't build them any other way. They are also likely to be very cheap because most any method for manufacturing things this small is to replicate them like crazy and make many at the same time.

Optical computing offers an interesting and orthogonal set of tricks and techniques. Note that in the cases I am describing here the invention that must be done is primarily engineering rather than scientific in nature. The basic phenomena exist in the lab today, so mother nature has, in effect, *already signed off on the designs*. The remaining steps are just to learn how to manufacture these devices in quantity, integrate them at a large scale and bring them to market. People are pretty good at doing this sort of thing, so I think that there is as much reason to believe that breakthroughs will accelerate the pace of development as there is to doubt that the trend will continue.

Assuming that the price/performance trend in computing *does* continue, the computers of 20 years from now will be a million times faster, and 40 years hence a trillion times faster than the fastest computers available today. In order to put this into perspective, a factor of one million reduces a year of computing time to just 30 seconds. A factor of a trillion takes *one million years* into the same 30 seconds. Attempting to extrapolate what we could do with a CPU year of 1993 level computing is hard, and a million CPU years is nearly impossible to imagine.

Note that this is only the estimate for a single CPU and the standard serial "von Neuman" architecture. Multiprocessors will become increasingly common, and give us a huge range of performance above and beyond this. The figures above were meant to be a comparison at fixed cost - say for a typical desktop computer. There will always be specialized problem areas in science, engineering and elsewhere that will justify machines that are 1000 times to 10,000 times more expensive than a desktop user of today can afford. The supercomputer of 2033 could easily be 4 quadrillion times faster than a computer of 1993 - which means that it could do in 30 seconds what the fastest PCs of today could accomplish in 4 billion years - the age of the earth.

Storage capacity is also increasing at an exponential rate. There are a whole host of interesting storage technologies on the horizon. Current hard disk technology will continue to drop in price - earlier this year we priced hard disks at about \$1 per megabyte. This fall a new series of disks will come in at about 30 cents per meg, and I believe that this trend will continue for a while. We have talked to people at disk companies who are planning to make gigabyte capacity drives down to the 1.8 inch form factor, and then surface mount hundreds of them to cards, as if they were chip modules, to increase the density. The total bandwidth will be very high because the disks will be configured as an array. Various forms of optical storage will also appear to challenge this standard magnetic based technology. The most exotic of these is holographic memories which appear to be able to store up to  $10^{12}$  bytes - a terabyte - *per cubic centimeter*. This isn't even close to the theoretical limit, which is far higher.

These forms of secondary storage will face some interesting competition - semiconductor RAM has increased in density (and decreased in cost) by 4X every 18 months for the last twenty years. The various technologies discussed above for CPUs will also effect mass storage - RAM based on single

electron transistors will be very dense. Various random factors (such as the recent explosion at a Japanese epoxy plant) can perturb the pricing, but in the long run I expect these pricing curves to be maintained. This is far faster than the price of mechanical mass storage is dropping, and suggests that by the year 2000, RAM will cost about \$1-\$2 per gigabyte.

I expect that the typical desktop PC around the turn of the century will be over 100 gigabytes - whether RAM or a mixture of RAM and some other sort of mass storage., and a typical LAN server here at Microsoft will have a few terabytes. To put this in perspective, the American Airlines SABRE reservation system is a little over 2 terabytes, so it would fit on only 20 PCs worth of storage. If you bought that capacity today with PC industry components the cost would be just under \$700,000 , a tiny fraction of what it cost to built SABRE out of 112 IBM 3390 disk drives. By the year 2000, I expect that amount of disk space to cost a few thousand dollars. In fact, it might cost even less because video on demand systems used to replace Blockbuster and other video rental stores will dramatically increase the market for storage and should dramatically drive the price learning curve.

The computational load used by SABRE is even less of a problem than the storage requirement. The existing reservation system uses processors (IBM 3090 architecture) rated at 423 million instructions per second. This is equivalent to four MIPS R4400s, or six to eight Pentiums. Another way to look at this is the traffic load. The peak ever recorded on SABRE was 3595 transactions per second. A dual processor MIPS R4400 machine over in the NT group recently benchmarked at about 300 transactions per second<sup>1</sup> with NT and SQL Server, so you'd need a dozen of these machines to do the whole thing. The reason for the factor of six difference between this and the raw number of compute cycles is that SABRE uses specialized transaction processing operating system (IBM Transaction Processing Facility) rather than a general purpose system like NT and SQL Server which is far less efficient (but more flexible) than TPF.

In fact, with a new generation of transaction and database software based on direct use of 64 bit addressing and massive RAM memory, a single microprocessor such as the MIPS T5, Intel P7 or next generation Alpha or PowerPC should be able to the entire SABRE load. Although SABRE represents one of the most challenging real time data processing tasks of 1993, hardware technology will pass it by in the next couple years. There is still a major task in creating the *software* which will enable this, but my assumption is that we and/or others will recognize this opportunity and rise this to challenge.

One way to consume these resources is to switch to ever richer data types. We have seen how the move to GUI consumed many more CPU cycles, RAM and disk space than character mode, so it is relatively safe to suggest that more of the same will occur with rich multimedia data types. Audio is already within our grasp, because it is fairly low bandwidth. Film and video represent a major step upwards. A feature film compressed with MPEG or similar technology will be about 4 gigabytes. Advances in compression will probably take this down to about 600 megabytes and increase the quality but this is still a major step up from text. If we compare the size of a novel or movie script to the film itself, we'll see roughly a factor of 1000 increase.

A factor of 1000 is quite a bit by normal standards, and much of the work that the PC industry will be doing in the next several years is gearing up to this challenge. It is quite clear that video and other multimedia data types will become as commonplace as text is today. This will cause major shifts in usage and will have all sorts of other consequences, some of which are treated elsewhere in this memo. Nevertheless, this isn't a very full answer to the question of what we'll do with all of those computing cycles.

Remember that a factor of 1000 improvement takes only *ten years*, so at best the transition to video will only be a temporary pause in the onslaught of computing. Other media types won't really help

much because there is a more fundamental limitation, which may be somewhat surprising - human sensory organs aren't all that complex!

Audio and video are already well understood, but we can make some estimates for the other senses. Taste and smell are actually variations of the same sense (you can't "taste" many things when your nose is plugged) which maps a set chemical sensors in the nose and mouth to a signal. Clinical tests similar to the "Pepsi Challenge" show that the taste/smell resolution is not very complex - at least in humans. The sometimes hilarious vocabulary used by oenophiles to describe the taste of a wine is another example. Without understanding the output device it is hard to say in detail, but I doubt that the combined bandwidth would be more than audio.

Touch is more interesting. Apart from our fingertips, lips and a few other regions, our sense of touch is actually not very good. The basic unit of visual displays is the pixel, short for picture element, and in a similar vein we can talk about the touch element or "touchel". Even parts of our anatomy which we think of as being quite sensitive often have poor spatial resolution and thus require few touchels to achieve the full effect. Poke yourself with various shapes of the same basic size and you'll find that most places you can't tell the difference - sensitivity as we normally think of is actually about resolution in *amplitude*, which effects the number of bits per touchel rather than their density. Even the highest resolution parts of your body don't need more than 100 touchels per linear inch, and the total surface area which requires that much is pretty small. An estimate of the total touchel bandwidth is probably something like one million touchels with 8 bits of amplitude per touchel, updated between 30 and 100 times a second. The total bandwidth is therefore equal to, or possibly a bit less than, that for video.

There are clearly some *major* engineering problems to be solved before we all have touchel based body suits and taste output devices in our mouths. Better yet we could have a direct connection to our nervous system so we can "jack in" to our PCs in the manner that William Gibson and others have described in cyberpunk science fiction. Since this is primarily a biological and mechanical engineering problem it is *not* subject to the same exponential technology rules which govern computing. There are reasons to be optimistic, but it is *not* clear to me whether we'll have the full range of human senses as output devices in twenty years, or even in forty. Regardless of the I/O issues, it *is* clear that computing resources will not be the bottleneck in solving this ultimate human interface problem. In fact, the estimates above suggest that the total bandwidth that can be absorbed by a human is only a modest constant factor larger than video, and thus will only take a few more years for the price/performance curve to surmount it.

Given the amazing growth of computing, there are always a foolish few who contend that we will run out of demand, so we'll never need more powerful machines. The counterexample is that there are many simple problems which would require more computing than a 1993, or even a 2033, level supercomputer could do between now and the recollapse or heat death of our the universe<sup>2</sup>. Computer scientists classify problems of this sort as "NP hard", and there are many examples. In actual practice some NP hard problems can be efficiently approximated, but many cannot, and this will lead to a never ending stream of problems for computers of the future, no matter how fast.

Here is the simplest example of why NP hard problems are hard. Consider trying to enumerate all possible combinations or orderings of N unique objects (characters or whatever). The number of combinations is N factorial, usually written as N!. 3! is only six, so I can list all of the combinations here 123, 132, 213, 231, 312, 321. 10! is about 3 million, which is relatively small, but 59! is about 10<sup>80</sup> and 100! is just under 10<sup>157</sup>. These are *really* big numbers! For reference sake, cosmologists usually estimate that there are about 10<sup>80</sup> elementary particles (protons, neutrons etc.), and about

$10^{160}$  photons in the entire universe, so even with extreme cleverness about how we stored the resulting list, we'd need to use *all* of the matter and most of the energy in the universe just to write down 59!, much less 100. As output problems go, touchel suits seem mundane by comparison! This particular example isn't very interesting, but there are plenty of NP hard problems, such as the traveling salesman problem and many others which are just as easy to state and just as hard to solve.

Many open problems in computing - including some classified as "artificial intelligence" fall into a grey area - they are likely to be NP hard in principle, but there may efficient approximations which allow us to get arbitrarily good results with bounded computing resources. Other problems, such as simulation, recognition based input (speech, handwriting...), virtual reality and most scientific programs are clearly *not* NP hard so they will fall to computing.

The key issue behind this point is that computing is on a very fast exponential growth curve. *Anything* which isn't exponential in growth, or which is exponential but with a slower growth rate, will quickly and inexorably be overwhelmed. NP hard problems are an example that can easily scale beyond any computer. Mainframe and minicomputers did have exponential growth, but at a slower rate than microprocessors so they succumbed. Human interface only scales up to the point where our nervous system is saturated, which as discussed above will be reached within the next decade or so.

The SABRE airline reservation system mentioned above is one which *doesn't* scale. The size of SABRE depends on the number of travelers, flights, travel agents and how fast they type. None of these factors is growing anywhere near as fast as computing itself. Even a constant factor of 1000 by using video, or of perhaps 10,000 to have a full virtual reality (thus obviating the flight altogether!) does not matter because *any* constant factor is rapidly absorbed by exponential growth.

This is a fundamental lesson for us. It is extraordinarily difficult for people to really grasp the power of exponential growth. No experience in our every day life prepares us for it. The numbers become so astronomically large<sup>3</sup> so quickly, as the projections above show, that it easy to either dismiss them outright, or mentally glaze over and become numb to their meaning. It is incredibly easy to fool oneself into thinking that you do understand it, but usually this just means that you've mentally done a linear extrapolation from the recent past. This works for a little while, but then rapidly becomes out of date.

The trick for the next decade of the computing industry will hinge on being very smart about recognizing which things scale with or faster than computing and what things do not. This will *not* be obvious in the early stages, but *that is where the value lies* so this is the challenge to which we must rise. The "MIPS to the moon" vision speech is very easy to make - *the hard part is to really believe it*. Fortunately, much of the competition will be a bunch of entrenched linear extrapolators that won't know what hit them.

At its essence, *this* is the secret of Microsoft and the personal computer industry. Despite the incredible advances that had occurred in mainframes and minicomputers, the people involved - many of them quite brilliant - could not grasp the fact that the advances would continue and that this would stretch their business models to the breaking point. If the mainframe folks had stopped to make the exponential extrapolation - *and acted upon it* - then it naturally follows that microprocessor based systems would deliver computing to the masses, that they would ultimately surpass mainframes and minis, that hardware should be decoupled from software because the driving forces are different, and finally that software would be a central locus of value.

The last point is quite interesting. If you are in the hardware business you must be very agile indeed

to keep your footing. The implementation technologies change so fast that you are always at an incredible risk of becoming obsolete. It is very hard to change a manufacturing business fast enough because of the large investment in tooling up for what will soon become yesterday's technology. The only hardware companies that have ever made significant money are those that managed to create an asset - the hardware architecture - which was above the fray of individual implementations and thus could enjoy a longer life span. Software is able to do the same trick in an even better fashion. Like a hardware architecture it lives for a long time - more than a dozen years so far for MS DOS - but the tooling cost is far lower. The software business is still quite tricky but it is fundamentally better suited to long term growth in an exponential market than hardware is.

The growth curve of Microsoft is unprecedented in the annals of business, and seems quite miraculous until you realize that what we have done is ride the exponential growth curve of computer price/performance. *That* is the true driving force behind our success. As long as computing hardware increases its price performance there is a proportionate opportunity for software to harness that power to do things for end users. This occurs because the machines get cheaper, user interface techniques like GUI allow us to make computing more appealing to a wider audience, raw power enables new application categories, old customers can upgrade and get substantially better features... All of these factors combine to increase the opportunity for software. I made a chart a couple of years ago plotting the number of lines of code in several of our products over time. These also followed an exponential curve with approximately the same growth rate.

Our feat, and believe me I do *not* mean to belittle it, has been to ride this wave of technology and maintain or increase our relative position. The correct way to measure this position, and thus our market share, is as a fraction of worldwide CPU cycles consumed by our products. (As an aside, the evolution of the information highway will cause this market share metric to evolve as well. Rather than measuring the fraction of the worlds computing cycles executed by our software, we will have to look at share of both CPU cycles as well as total data transmitted.)

Maintaining CPU cycle share is very hard to do, because the right mix of products and technology to do this at various points in time changes a *lot*. Companies that have point products or which only extrapolate linearly will fail. So far, we have often been the beneficiary because we have been able to reinvent the company at each point along the line.

The bad news about this is that we will have to continue to do so, and at a rate which continues to increase. It will impossible for us to maintain our historical growth curves unless we maintain our CPU cycle share and that means that we must do two things: continue to bring new technology to our existing products, and at the same time create new product lines to track the emergence of computing in new mass markets.

One aspect of the price/performance trend discussed above is that a PC class machine will get amazingly powerful, but an equal consequence is that extremely cheap consumer computing devices will emerge with the same or higher computing power than today's PCs, but with *far* higher volume. The bulk of the worlds' computing cycles come near the low end, high volume part of the market. Any software company that wants to maintain its relative share of total CPU cycles *must* have products that are relevant to the high volume segment of the market. If you don't, then you are vulnerable to a software company that does establishes a position there and then rides the technology curve up to the mainstream. This is what the PC industry did to mainframes and minicomputers, and if we in the PC industry are not careful this fate will befall us as well.

The basic dynamics of this situation comes from the relative lifetimes of software and hardware. An operating system architecture can easily last 15 to 20 years or possibly even more, as MS DOS and UNIX both demonstrate. Over that period of time the computing hardware that the OS runs on will

increase in price/performance by a factor of 30,000 to a million, taking the combined platform into new realms and application areas. If each segment of the computing industry scaled at the same rate, then the relative standings would be safe, but this isn't the way it usually works. Technology trends will favor some segments over others. Microprocessors beat out discrete components, and more recently RISC based microprocessors have been increasing their price/performance at a faster rate than their CISC based rivals. Business trends also matter - a high volume platform will have more ISVs, and probably a more competitive hardware market.

Microsoft started out as the Basic company - well, as it so happens we *still* are the leading supplier of Basic, but that is hardly the way to characterize us at this point. Growth, profitability and the focus of the company have historically shifted from one area to another, and this will continue past our current product lines. The day will come when people will say "hey, didn't Microsoft used be the company that made office software?" and the answer will be "yes, and as a matter of fact they still do, but that isn't what they're known for these days".

## Your Life On Line

Consider spoken conversation. Although we may quote each other when we say something memorable, that is the exception rather than the rule. People tend to treat speech as an ephemeral medium - if we want to count on being able to revisit it in the future we go to special effort to take notes or record it. In the near future we will have the capability to record and index it all.

It's pretty easy to compress human voice down to about 4800 bits/second today, and this will almost certainly improve to 2400 bits/second in the near future, which works out to about 1 megabyte per hour. The DAT drive that I use for backing up my hard disk stores 5 gigabytes on a single tape which works out to 5000 hours. If you assume 8 hours of sleep a day, there are 5824 hours you're awake per year. Of course there is no reason to record the periods of silence, only periods where you are speaking or somebody is speaking to you. Thus my single DAT tape can probably hold many years worth of audio. It's actually an older unit - the next generation of digital VCRs coming out in 1994 will hold 100 gigabytes per tape, which means that a single piece of media with a cost of about \$10 should be able to hold all of the conversations an individual has for their entire lifetime.

Video takes considerably more room than audio. If you assume that this sort of archival video can be a bit lower quality than standard VHS, it will take about 1.5 megabits per second with off the shelf compression methods. This means that recording video onto digital would cost about *eight cents per hour* using digital video tape, or a couple hundred dollars per year. Note that this is a 1994 cost - within a few years it will be far cheaper.

Location is another important attribute. This can be obtained with a small GPS receiver, or a variety of other small wireless devices. Recording the precise location (within a couple of feet) and the precise time of every movement a person or vehicle makes is far smaller amount of storage than voice. While we are at it, we could record temperature (ambient as well as body temperature), barometric pressure, blood pressure and a variety of other data about you or your surroundings.

Those of us who work with computer will also create files, receive email and so forth. A really fast typist can do about 100 words per minute with a 5 character word length, or about 30K bytes per hour. Lossless compression can typically reduce this to about a third, so this is about 100 times smaller than voice. Mouse motions and application commands are tiny on top of this, so we certainly have the capability to spool all keystrokes, commands and files that an individual creates

over their lifetime at a fraction of the cost of voice.

The result is that it is easily possible to create general digital recorders which will spool audio, video and other data very cheaply. Cryptographic techniques can be used to time stamp or location stamp these records with digital signatures so that they cannot be forged or tampered with. The entire input and output of an individual's life can be spooled and recorded onto digital media. From there it is a small matter to index it for retrieval based upon time, location and content (via off line speech recognition).

The first application of this will come in the various individual systems. Given the increase in storage on PCs, why not record every version of every file? High speed networks and new software will make this quite cheap. Typing and other input mechanisms don't scale with the growth of computing, so why not spool every keystroke and mouse move. The notion of losing work should go away, because it is trivial to record *everything* that a user does. Caller ID telephones have started to do this with a list of every caller (answered or not) - why not record the content as well?

The Rodney King case showed the power of video tape in recording a situation. How long will it be before every police car, or individual policeman, is equipped with a digital video camera, with non forgeable time and location stamps? Commercial airliners have long had a "black box" or flight recorder, but why not extend this for buses, cars - even for individual people? In some cases this would be limited to time, location, speed and in others it would record full video and audio.

In fact, while we are at it, a unit like that is cheap enough to be deployed even more widely. Banks and some stores have cameras which can record in the event of a crime, but given the low cost why not have them nearly everywhere? Every city street light represents an investment by the community toward public safety, yet these lights are far more expensive to install and operate than a digital video black box will be in a few years. I could easily imagine such automatic recorders positioned in literally millions of places; connected to storage and analysis facilities by the information highway.

A couple more years of price/performance improvement and we could carry a device like that with us. The digital storage capacity would be enough to record every thing we say and do - for our entire lives. Such a device could carry your entire life on it - the ultimate diary and autobiography.

At this stage you may wonder why do this? Indeed, the notion of never being outside the range of digital recorders operated by Big Brother can sound pretty sinister. It would be very easy to regard this as the ultimate threat to privacy and personal freedom. At another level it may be pretty stupid. Everybody has, at some point or another, told a lie or done something that in retrospect they aren't proud of - why should we record ourselves and make this obvious? Richard Nixon's decision to record conversations in the White House for future historians clearly backfired and he was hoist on his own petard.

The social, legal and moral implications are quite interesting because our existing traditions are not prepared to cope with the incredible potential to record and analyze information on this scale.

There will be many arguments about how to extend the principles we have today to this new world. In fact, we have already started to see the first skirmishes in this battle in the FBI proposal to make encryption illegal. Their proposals are chilling to a technologist, but their argument is disarmingly simple. They just want to preserve the abilities that they have today. Sounds reasonable, right?

Regardless as to whether putting your life on line is good or bad, it is very clear that it will be both feasible and quite cheap. Given this I believe that it will be widely used in at least some circumstances. Here are some scenarios as to how and why we may find our lives on line in the future.

The first is that there are many contexts in which it would be very useful to have a personal record of

what you've done. This is clearly true for computer files, where returning to an old version is already done quite often. Researchers at EuroPARC (the Cambridge branch of Xerox PARC) have created a system that uses wireless electronic badges to record all of the motions of people inside a building, then create a "biography" which says where you were at any point and who you were with talked to (they do not record voice yet).

Unfortunately there are other reasons besides simply utility to the end user. Insurance companies would love to have the black box recorders on your car - both to trace it if stolen (there are several systems, such as Lojack that do this today), and to provide information in accidents. It isn't hard to imagine having half price car insurance if you have a car with a recorder in it. This can quickly extend to other circumstances. Medical malpractice insurance might be cheaper (or only available) to doctors who record surgical procedures or even office visits.

Many employers will want to use these recorders as a management tool. Bus and trucking companies have an obvious interest in whether their drivers stay within the law and are on time. The public at large may insist that the police record themselves in the course of their work. This guards against claims of brutality or abuse on one hand, and gathering better evidence on the other. Other circumstances will be more invasive and harder to justify, but it is virtually certain that people will try.

Personal safety will be a big issue as well. There are plenty of situations where you'd rather have Big Brother watching than be left to the mercy of street crime. David Brin's science fiction novel *Earth*, suggested that digital video cameras with wireless links to the police would revolutionize police work and the judicial system. In addition to putting them in street lamps, Brin has them toted by senior citizens, because retired baby boomers will make up a large fraction of society. Since they will also dominate the voting population, the laws of his imagined society are oriented toward widespread use of digital surveillance to protect senior citizens against crime.

One of the biggest reasons will be defensive, using the digital record to protect against litigation and other claims. People who record are open to the possibility of the recording being used against them, but at the same time will be able to protect themselves against spurious claims. If you have a business deal with another party, and they have selectively recorded things (not every meeting etc.) then the best defense is to have a more complete recording yourself. One side effect of cheap computing is that it will be possible to forge essentially any kind of photograph, video or audio. The only such data that will be admissible in court will be those taken with tamper proof hardware that digitally signs it.

If these trends continue it will be gradually become more and more accepted to have these records - both at a corporate and personal level. What today seems like the digital jackboot of Big Brother will one day become the norm. Every year the executive staff at Microsoft have to sign a "boy scout pledge" of business practices which says (in essence) that we see no evil, hear no evil and do no evil. I'm not trying to make fun of this - it is true that our policy is to be good guys - but it doesn't really change from one year to the next. The ritual of signing this annually small testimony to the societal pressure to redundantly reaffirm our values. Twenty years hence the pendulum of public opinion may have swung to the point that people say "What do you mean, you don't record everything? Do you have something to hide?".

## **Gutenberg Reprised**

If you grant that the world writes and makes decisions with PCs, what is next? The real answer is long and complex, but three of the key components are to *read*, *communicate* and be *entertained*.

An even simpler way to describe this is to say that computing technology will become central to *distributing* information.

As a rule, distribution has much more pervasive effects than authoring. Improving life for the author of a document does not materially effect the size or nature of the audience that she can address, but changes in distribution have a dramatic effect. The clearest precedent is the invention of the printing press. Great works of science and literature - Euclid's geometry, Plato and Horace, the Bhagavad-Gita, the Iliad and Icelandic Sagas, all- existed long before the printing press, so humans clearly were able to conceive them, but they had a very limited customer base. Monks and scribes spent lifetimes copying books by hand, while bards and minstrels memorized and orally repeated tales to spread and preserve them. No matter how cheaply one values their time, it was still a very expensive proposition which was the primary limiting factor in broadening the number of customers. If we could use a time machine to supply all those monks with PCs and Word for Windows, but limited the rate at which they could print to the same level of time and expense, it would make little difference - except perhaps for letting the monks channel their energies toward other fields<sup>4</sup>.

When Gutenberg did change the economics of distribution, the world changed in a fundamental way. It is estimated that Europe had on the order of ten thousand books just prior to Johan's invention - within fifty years it would have over eight million. Literacy became a key skill. The advent of mass media - through printed handbills - revolutionized politics, religion, science and literature and most other factes of intellectual life.

I believe that we are on the brink of a revolution of similar magnitude. This will be driven by two technologies - computing and digital networking. We've already discussed the change in computing technology, and that is certainly dramatic, but it is communications which really enables distribution.

The technological factors which will drive communication are very simple. The first is that wide area communication will be digital and will therefore be driven by the *same* price/performance curve of computing as a whole. Switches and other parts of the new digital communications systems will be computers. Today this is partially true, but most of the world's telephone system is still analog, and the digital components that are in place are constrained by the old analog architecture (circuit switching etc.). If you made a relative comparison to the computer industry, you'd find that a "state of the art" central office switch of today, like the SS7, is comparable to a mainframe or minicomputer of the 1970s. In fact, the SS7 is built out of a couple of AT&T 3B2 minicomputers! As the microprocessor revolution sweeps this industry we'll see computer style similar changes in the price/performance of the switches and related equipment.

That covers the cost for the boxes at the ends of the line, but what about the stuff in the middle? Fiber optics driven by semiconductor lasers will be the primary connection in between the switches and terminal equipment. Fortunately, this technology is cheap, powerful and scalable. The basic fiber itself is just drawn glass which is able to channel a beam of light. The magic comes in when you can efficiently modulate the light with a signal, demodulate it on the other end, and then replicate the whole thing on as many strands of fiber as you like.

The result of these two technologies will drive the cost of bandwidth per customer down by enormous factors - in fact by about the same amount as the amazing numbers discussed above for computing. This is the key to changing the economics of distribution, and it is central to the entire notion of the "information highway". I'll use this general term to mean wide area digital consumer networks. In most cases I'll be talking about high speed networks which can transmit high quality video or other digital data. In the long run (actually not all that long!) all wired communication will be broadband. Nonetheless, we should not forget low to medium speed networks (the existing phone

system or ISDN) because many of the essential features of the information highway are already exhibited by these systems and for some data and services they are quite suitable.

## The Tyranny of Geography

One of the most general and dramatic aspects of the information highway is to virtualize space and time. Put another way, the highway will break the *tyranny of geography* - the stranglehold of location, access and transportation that has governed human societies from their inception.

An old maxim in the real estate business is that the three top factors that determine a properties value are *location, location and location*. The reason is obvious - physical location and access is an incredibly important thing to a commercial building or home. It determines how many people come in the door of your business, how far people have to commute, who your neighbors are and a myriad of other attributes.

Similar factors govern much of society. Transportation systems have been incredibly influential to the development of cities, towns and suburbs. Every major change in transportation has caused enormous changes in our lives, and in the business community at large. Access to people is the primary differentiator between city and suburb. Most of the differences between a street in New York City and one in Redmond can be traced to the traffic flow and access potential- how many people can a business on that street draw its customers from and how far must they travel? How many jobs are within commuting distance of a home on that street?

Any activity which can be transacted over the information highway is freed from these constraints. Assuming that the network is there, a person in Redmond or Manhattan or nearly anywhere else will have *equal* access to goods and services presented on the network. Geography doesn't matter any more - if you are on the net you are in a virtual world which is not bound by these conventions or constraints.

The implications are enormous - so much so that it is hard to overstate them without sounding silly. *Any information based business or activity will cease to be geographically bound*, which means that the local operators of these activities will die. Nationwide services which are based on the information highway will dominate them every time.

This will occur in various stages over the next couple of decades. The first examples will be pure information businesses, which can simply transfer digital data directly over the network. Software (and therefore Egghead and others) will certainly go this way, as will the local video rental store. The next step up is information used to purchase or deal with physical products - mail order and retailing. This will have a huge effect on local retailing, although it will not eliminate it.

Much of the focus on information business is placed on mass market retail information, but there is another category that is ultimately even more important - wholesale point to point information - in other words, the work product of most people employed in offices. Telecommuting will allow these people to go to work over the information highway rather than the physical one, and thus be bound only by its constraints. There is still a speed limit on the information highway - but it's 186,000 miles per second<sup>5</sup>. Long distance bandwidth will drop in price until there is very little reason for people working "together" to be nearly anywhere on earth which has network access.

Initially this will just effect a fairly small number of people, but over time this could become an enormously important phenomenon. An industrial economy needs to have its workers near their factories, but our increasingly service and information oriented economy can let its people disperse

to wherever the information highway can reach them. This will lead to de-urbanization of the developed world. Most of the folks who work in large office buildings in Manhattan and commute in from Connecticut or Long Island would be a lot better off telecommuting - at least part of the time. Ultimately, "suburbs" will not be constrained to be near urban areas - any area with decent network access will do.

This will take both time and technology to occur. People's customs do not shift overnight, and we will need to come up with lots of tools and support. Video conferencing applications, clever user interfaces and a variety of other things must come into play before people work via the highway on a wide scale basis. Even though the pace of change may be slow, the total effect will be very large.

Perhaps the ultimate expression of geographical tyranny is found in politics. My vote is aggregated with those of my neighbors - first to vote for local government, then state offices, the House and Senate and ultimately the President of the United States. At every level, geography plays a huge role - it determines who I can vote for and whether my vote will be heard amid the din of the others who are aggregated with me. In fact, political control of the aggregation process - known as "gerrymandering" - is widely used as a mechanism to manipulate the outcome of elections.

The tacit assumption in a geographically based representative democracy is that I have something in common with my neighbors. Since we have this common interest, we should pull together to elect representatives to carry our message - either to the state house or the nation's capital. This may have been relevant years ago, but it is certainly fallacious today. I don't have anything in common with my neighbors on most issues that are important to me. There are some legitimate things that are truly based on local geography, like whether to widen our street, but on the really important topics I probably have far more in common with a computer industry person in Massachusetts or Silicon Valley than I do with the retired couple that lives next door to me in Bellevue.

This trend has already become moderately important in politics. Causes like environmentalism (e.g. the Green parties in Europe), gender or sex (gay rights, feminism...), moral issues (abortion) and others are strong enough to polarize people and create movements that cross geographical boundaries. The limiting factor is that it requires a lot of work to form and manage these grass roots organizations and get the word out.

The information highway will greatly accelerate this trend because it enables virtual communities and new forms of communication. It is very expensive for one person to communicate with more than one other person - the telephone network is fundamentally a point to point service, while radio and TV go to the other extreme and force broadcast communication to everybody. Bulletin boards - initially text and later video and multimedia - are an effective and inexpensive way for groups large and small to engage in multicast communication. Every interest group will have its bulletin boards (text, video...) its video productions and its ability to conduct polls and plebiscites on the network. The information highway will become a powerful conduit for grass roots political organizations which will eventually exceed the power of the traditional media.

Once the populace is thoroughly ensconced in its BBS forums and interest groups, there is still the problem that aggregation is done in the archaic geographical districts. As the power of the highway in reaching and influencing people grows, there will be a clamor to redress this situation. Television changed politics from the cloistered world of the smoke filled room where party officials made all the key decisions to the world of the sound bite and the infomercials, because it is TV that gets out the vote these days rather than ward heelers and party bosses. New politicians<sup>6</sup> and causes will emerge to exploit the information highway's ability to reach and motivate people, and it will become the most leveraged way to get out the vote.

The tension between the old geographically based politics and the new highway based approach will ultimately have to be resolved. One solution might be will be to elect representatives from interest group constituencies rather than districts. The simplest way to do this is to implement Lani Gunier's ill fated proposal to use cumulative voting to empower minority groups<sup>7</sup>. This was originally intended with ethnic minority groups in mind, but in practice it could reach beyond this. The true minority groups are not those based purely on race or ethnicity but rather on the specific issues and values that smaller groups of like minded people share, regardless of race or creed. If each American got 535 votes to cast any way they like against a national slate of representatives for Congress there would be a very different set of people in the House, and a very different set of political dynamics.

The next logical step is to make government itself use the ability to communicate and vote on issues. The information highway can conduct votes at nearly zero cost and with far greater assurance against miscounting or fraud. Representative democracy is based on the notion that the people cannot feasibly be consulted on every issue, but the information highway fundamentally enables us to do exactly that.

Note that I'm *not* claiming that any of these changes will be better than the current system (although it is hard to imagine that it could be any worse!) The point in politics is rarely about objective good versus bad - it is about what is popular. Ross Perot made "electronic town halls" part of his platform, and it is entirely possible that he or someone else will impose these changes in a top down fashion. A more likely scenario is that the information highway will become a vehicle for traditional political discourse first, and then slowly change the system itself as people become increasingly frustrated with the tyranny of geography.

## **Sam Walton On Steroids**

WalMart is one of the great success stories in retailing. The basic strategy behind it was very simple - move business away from specialty stores, particularly in small towns, and towards large warehouses which can have lower overhead and therefore lower prices. It was a great idea for everybody except small town merchants, who have in effect become kindly old Sam's personal road kill.

Dell Computer, Gateway 2000, Circuit City, CompUSA and their ilk have done something roughly similar to old style computer stores, and a good portion of direct sales force activities. Whether you are a mail order firm or a superstore warehouse is only a small distinction - the key issue is maximizing efficiency in physical distribution.

The interesting things about these trends is that they will both become far stronger as the information highway develops - mail order is made easier, and the warehouse store doesn't have to be open to customers because they can browse it on the highway. The effect on retailing will be very large, as discussed in a section below, but there is another more general point that I think is worth making - many of the most powerful effects of the information highway, especially in the early days, will come by *accelerating trends that are already well under way*.

The reason for this is pretty obvious - if a trend is already developing and it is compatible with the highway then it will be happen faster. Human inertia is difficult to overcome and that will be a big barrier to things that are truly novel and require major changes in behavior. A change which is already in the works and has the groundwork underway, or which has already been proven to be an

easy adjustment for people, has much less of an inertia barrier than things which are totally new. This point may seem so obvious that it hardly bears repeating, but it does not appear to be well understood. Time and time again I run into people who make statements about what will or will not happen on the information highway which are totally ludicrous when you consider that the same logic would contradict things which are already important in society. The converse is that reasoning by analogy with current trends is one of the most important analytical tools for recognizing road kill in the making.

## **I Brake For Couch Potatoes**

Do people want to interact with their TV sets? Aren't we really a nation of couch potatoes who'd rather vegetate than explore the information highway? How much change to people really want? Is the couch potato going to become road kill?

These are fascinating questions because it is hard to accurately judge the amount of inertia in human behavior patterns. It is certainly within the realm of possibility that people will be unable or unwilling to accept the changes that come with the information highway, materially slowing its growth.

The couch potato versus interactive debate is usually portrayed as a black and white issue - either people vegetate in front of their sets, or they are dynamically interacting hitting a button a minute. This characterization serves the purposes of the various zealots and luddites who pontificate on this topic. Regardless of whether you are for or against interactivity, it is easier to make your case if you assume a polarized either-or situation.

In actual fact, I think that it is far more likely that we'll find a very inhomogeneous distribution of interactivity - both demographically and by application or service. Some people are confirmed couch potatoes today and they may remain that way indefinitely in some modes of operation, but I find it hard to believe that there will be a large number of people who love the status quo so much that they won't take advantage of interactivity in at least some way. Even the most sedentary couch potato uses the telephone, pours through classified ads, sends greeting cards, looks things up in the yellow pages, calls in to order take out pizza, purchases things advertised on TV and other activities which will be available as interactive services on the information highway. You'd have to be in an intensive care ward to not be interactive in some aspects of your life.

Those who claim that most people are couch potatoes are usually taking a far too narrow definition of what the information highway is about. Even if movies and TV series - primary fodder for the couch potato - do not become interactive there will still be many other things that are. The real question is not whether we all start blasting away interactively whenever we sit in front of a terminal for the information highway, but rather whether enough services exist to fund the continued deployment of the highway. This doesn't take a great deal of interactivity in the early days.

The couch potato debate is really a very narrow issue about the nature of entertainment - will large numbers of grown people use entertainment formats which are active or passive? The answer in the short term will depend a lot on the details - will creative people come up with compelling interactive titles and services? It is hard to predict whether this will be terrific games, or daytime soap operas with interactive plot control or a million other possibilities. There will be both ample potential and incentive for people to try and come up with these new forms and formats.

In the long term I think that they will succeed in this quest for two reasons. The first is that people

will have time to experiment and refine the medium. *The Milton Berle Show* is pretty odd by current television standards, and *MTV* would seem pretty strange if you could take it back in time to early television audiences. It may take 10 or 20 years before the information highway really hits its stride (television certainly did) as a mechanism for entertainment, but in the grand scheme of things that is a perfectly acceptable solution.

The second reason is a demographic one. We already know the power of Nintendo and other video games for children. One out of three American homes has one of these devices and their children spend a lot of time with them. We are raising a generation of people who have grown up with interactivity and *love it*. The information highway will be a far smaller change for these people than for those of us who only experienced computing as adults and are already set in our ways. As the Nintendo generation grows up and enters the economy the information highway will have a powerful set of advocates.

## **The Virtual Proletariat**

The gaps between rich and poor, whether at an individual level within our society, or at a global level between rich and poor nations have up to this point been based on a broad set of factors, both social and economic. The advent of the information highway poses a new kind of gap - that between those who are "wired" and have access to the highway and those who do not.

The new "haves" who are connected to the highway will, over time, will command an enormously greater range of opportunity than the have nots who aren't on line. In the early stages there really isn't much of a social issues because the very first systems will be dominated by video on demand and other forms of entertainment which hardly count as a serious difference. As commerce, education and broad scale communication move onto the highway the difference will become quite material to an individuals ability to be part of mainstream society. This isn't something sinister - it is merely a reflection of the fact that the information highway will be an incredibly powerful thing, and directly in proportion to its power it will disenfranchise those who do not have access.

On the other side of the coin is an equally powerful phenomenon - once you are on the highway there is fully egalitarian access to on line resources. All users are created equally - both geographically and socioeconomically. There is no technical or business reason why a person in Beverly Hills or Medina should have different access than somebody in Boise, Bozeman or the South Bronx. Some services may cost more than others, but fundamental access to communications, reference materials and other important aspects of the information highway can be available to essentially anybody. Virtual equity is far easier to achieve than physical equity, because it is driven by the exponentially exploding capabilities of computing and communication. The information highway will not by itself eliminate the barriers of prejudice or inequity, but in the long it has the potential to be a powerful force in that direction.

## **Markets Meet Their Makers**

The market place is one of the central concepts in both commerce and economics, and occupies a central place in any human civilization. The modern incarnations of the open market place are usually quite specialized, but at essence there is not that much separating the New York Stock Exchange, the Chicago Board of Trade, Sotheby's, a North African bazaar in Tangiers or the Tsukiji

fish market in Tokyo where sushi chefs go to buy the day's catch. In each case buyers and sellers gather in a location, shout and argue about price, and ultimately exchange goods.

The existence of a viable marketplace is an enormous boon to commerce and trade. Without stock markets it would be very difficult for companies to raise capital, and without produce markets it would be impossible for farmers to sell fresh fruit and vegetables. The overhead associated with making sure up front that there was a buyer for each tomato would be an enormous burden. An open market generally means competitive pricing and a reasonable degree of efficiency - goods move from buyer to seller with a fairly modest "friction" due to the people who operate the market - the "market makers".

In the case of Sotheby's or a similar establishment, this is the auctioneer. In the case of the NYSE, it is a rather bizarre chain of people starting with your stock broker, moving down to traders who have "seats" on the exchange and then finally to a "specialist" who manages the market of one or more specific stocks and matches buyers with sellers. That is for one side of the transaction - a similar set exists for the other party in the transaction. All of these middlemen take a piece of the action. Gambling with a bookmaker is actually quite similar - the bookie is the equivalent of the specialist in that he matches bets on both sides of horse race or sporting event, with just enough difference between the two to take a piece of the action, known as the vigorish or "vig".

Despite the tremendous success of markets such as the NYSE in raising capital for corporations, the entire system seems quite quaint when compared to what the information highway can offer. Why are a lot of humans in the middle taking a vig? Why do they have to congregate in one physical location in one city? Why is the vig so large?

The information highway will become the ultimate market maker for goods and services - posting a buy or sell order is a far more direct way to mediate a market than any of the existing markets. Distributed databases can post offers, resolve offers into completed transactions, handle authentication and security and handle all aspects of the marketplace. No humans need be involved except for the buyer and seller. Specific markets on the highway may still charge a vig, but in general it will be smaller than anything charged today.

There is already substantial precedent for this. NASDAQ already operates as a distributed system, albeit a rather crude one, where various brokers distribute the market making function. A network of institutional investors with PCs and modems forms the basis for the Wunsch Auction System (recently renamed the Arizona Stock Exchange) which bypasses existing security markets for a smaller vig. As systems like this catch on and use the facilities of the highway to provide better services, the existing market makers will be unable to compete. Some of their services will still be needed - for example advice on which stocks to buy or when to buy them, or expert appraisals of art works, but these will be decoupled from the mechanism of the market.

The really dramatic step comes when brand new markets are created on the highway, so that goods or services which never had the benefit of widespread trading can be efficiently exchanged. Used cars are auctioned today, but to used car dealers rather than the public. There is no reason to make this restriction on the highway. Many classified ads attempt to establish an auction for a specific item when they say "best offer" for the price, yet the ability to really attract the best offer is limited.

In other cases, new "trading vehicles" will be created in order to turn a non-tradable item into something which can be bought and sold over the network. In the 1980s this occurred in a number of financial markets. A portfolio of many home mortgages was bound together, then shares in the aggregated pool were sold as mortgage bonds. This allowed banks and financial institutions to trade home mortgages, but without having to deal with them one a time. Futures options were

created on a variety of synthetic securities - such as interest rate futures, or stock market index futures. In a very different domain, time share condominiums created a trading vehicle whereby people could purchase a resort home at a different level of granularity - by the week rather than forever. Each time this occurred it generated massive amounts of volume and dramatically changed the businesses involved. The information highway will provide such a strong foundation for creating marketplaces that it motivate the creating of many new trading vehicles.

Within 10 years, the information highway will host the most efficient markets ever created for a wide variety of items, ranging from sophisticated new securities to household commodities. A wide range of today's market makers and middlemen will be left looking for something else to do.

## **Species of Road Kill**

The remainder of this memo will offer some vignettes on how the highway will effect - or flatten - a number of existing businesses and customs. This analysis has the usual caveats that pertain to any such extrapolation. I could be very wrong, but it is interesting food for thought nonetheless.

One very important thing to keep in mind is that when I talk about companies or industries becoming "road kill" it doesn't mean that they will go out of business tomorrow. The actual process by which the information highway will displace current businesses will be far more like the start of an ice age rather than an instant calamity. In the long run the changes will be massive, and in retrospect it may seem to have happened overnight, but from this vantage point it will be slow transition. The only dramatic part is that the window of opportunity for a company getting on another strategy and avoiding extinction may come and go many years before it becomes obvious that they are doomed.

## **Communications Equipment**

The market for communications equipment has been very lucrative for the last couple of decades, but the hardware trends discussed above will make it extremely difficult for this to continue. The hardware components that comprise switches, line cards, terminal equipment and forth will be driven down the price/performance curve by the trends in computing. In fact, the communications equipment industry is in a very similar position to that of the computer industry in 1981 when the IBM PC was introduced. The PC design became a public standard (quite unintentionally on IBM's part) which allowed lots of new entrants to build a machine which was compatible with an industry standards. The ATM standard for packet switching is the equivalent standard for communications equipment.

ATM switches, interface cards and cabling will fuel the growth of lots of start ups as well as new product introductions by major manufacturers. ATM will become the key standard at every level of the industry, replacing the LAN and PBX as well as the WAN and central office.

It is certainly possible that the major suppliers of PBXs and central office equipment could downsize their current projects and simultaneously ramp up agile organizations to make cheap ATM equipment; all the while learning to live on dramatically lower margins. Yes it is *possible*, but it is an enormously difficult task, and I expect that very few will manage to pull it off. Wang, Unisys, IBM, and DEC didn't manage to do this smoothly, and although one or two of them may rally they have all gone through enormously painful travails in the process.

This covers the hardware, but there is also a strong software component to communications equipment at both the PBX and central office level. During an interview last fall a reporter for the *Wall Street Journal* pointed out to me that AT&T had almost 10 times the number of software developers that Microsoft does. Despite this resource inequity I don't believe that the major equipment companies are in *any* better position to write the software for the new generation of ATM based communications systems than the hordes of IBM mainframe software developers were in writing PC software.

The individuals involved are probably smart and their experience may be useful - but only if they are taken out of the old environment and have their goals and philosophy reset. Many *people* made the transition from working on big computers to working on small ones, but in my experience no organization ever did.

In fact, software will be the area in which these new communications systems are hardest hit. Today's systems - whether for a PBX or central office, are notoriously hard to modify and are all proprietary - much like minicomputer or mainframe operating systems. There are few if any third party software developers for them, and direct programming by end users (or their MIS organizations) is impossible. As a result there are few applications and little flexibility.

The trend for the future will be towards third party applications which write to industry standard APIs - just like on PCs. Many of these apps will promote *end user programmability and customization* - both at simple levels like call forwarding, out of office and related features to advanced systems for use in vertical markets. The TAPI spec is the beginning of this trend because it allows Windows apps to interface to a PBX but many more opportunities exist. I believe that one of the next "killer apps" in the office (and later the home) will be communications control. Today the most that we get to do is set a few speed dial numbers, and record our answering machine or voicemail greeting. Programming your communications - including voice, video and conference versions of both as well as email and FAX, will become important to both organizations and individuals. As this occurs, the old style PBX and central office switches will be unable to compete.

The software architecture for supporting this new flexible approach to communications is likely to move to software companies rather than be written by hardware vendors. The hardware itself will be priced on the quality of the implementation and many different hardware manufacturers will support the same software. In other words, the situation will be a lot like the PC market. This is going to kill most of the current set of communications companies, and I expect that within five to ten years most of the important players will be new entrants similar to Compaq, Dell and Gateway, with few if any of the current leaders left as real contenders.

### **Novell ???**

People who make local area network equipment aren't going to fare much better as a rule than the folks in the communications business. The one saving grace is that most of these people already live (or already suffer) in a very low margin, fast paced market so the effects will not be anywhere near as traumatic.

There is a persistent hope on the part of many people involved with LANs that ATM and fiber networks will basically be a transparent change to their current business. This theory holds that ATM will primarily be used as a high speed version of Ethernet - probably for network backbones and other places where the throughput is required. There is no change needed, they say, because this is just a better version of what we already have.

I believe that this is ridiculous. The argument is tantamount to saying that PCs wouldn't effect the mainframe market because they'd be used as terminal emulators. While it is true that some PCs were used this way, particularly in the early days, the really dramatic and important aspects of the personal computing revolution were about *totally new applications* which fueled customer demand and opened vistas that were never part of the mainframe world. The same will be true for fast ATM based LANs - they will allow video and voice telephony, communication of rich multimedia documents in real time and seamless integration with high speed wide area networks in ways which current LANs are wholly incapable of providing. In doing so, they will create a new set of users and applications.

The biggest casualty in this transition will be existing network software architectures, because they too are wholly incapable of serving the new world. There are a couple of fundamental reasons for this - first off, they have no ability to deal with real time, yet this is an essential aspect of delivering video and other multimedia streams.

The second reason is bandwidth - the basic protocols and APIs involved will dramatically limit the network performance to a tiny fraction of the network hardware potential. Remember that these new networks will be *faster than the local bus* on current PCs! It is no more appropriate to use NETBEUI or the Novell equivalent to access these networks than it would be to use them to access the video display or between the processor and RAM. The whole role of network protocols shifts because for the most part the protocols must *get out of the way* and let the data flow directly. This requires a very different architectural approach to networking.

The third and final reason is that for the most part the current LAN software architectures do not address the key problems that will arise in ATM based LANs that are interconnected with ATM based wide area networks - a lot of new functionality will be required. The basic things that people use Netware for today are really pretty boring - file and print service and email. These needs will continue, and can be hosted on top of an ATM network, but I doubt that they will be the primary driving force. Instead, the key asset will be a new generation of operating system and network software specifically designed for both LAN and WAN use.

This is an asset that we must try to secure for ourselves. Novell is also in a position to move in this direction, but in doing so they don't have any advantage over us because the departure from their current world is so great. With luck, they won't figure this out (we certainly should not tell them!). There is some evidence of that they have at least an inkling because they bought Fluent, a multimedia networking company, and they appear to be sniffing around the cable industry.

Even if they do move in this direction (and they already be) there is every chance that we can beat them because we have a variety of projects (our cable deal, our narrowband platform, AtWork...) that have far more synergy with the radical approach than they do. If we are successful at this, then Netware could become passé, and Novell could wind up as the biggest piece of road kill on the information highway.

Of course, if we allow them to win, then they will have renewed and extended their franchise in a very dramatic manner, and our hopes of unseating them will be just as dead.

## **Communications Carriers**

People who provide communications services face a somewhat similar problem to that of the vendors who supply them with equipment - the dramatic price/performance changes that come with computing equipment.

The direct cost of equipment isn't the only factor in pricing communications as a service - there is a lot of infrastructure cost which depends on things like the cost of having guys with hard hats and a backhoe dig up the street and lay a fiber optics cable. Regulatory issues also change the costs and create barriers to entry that are independent of the technology. Nevertheless, the advent of new technology will cause huge problems for existing carriers, because despite the cost and regulatory barriers, it is *still* true that any existing communications infrastructure will be at a huge disadvantage compared to a new one based on the latest technology.

The clearest example is the race for the "full service network". Cable TV and phone companies each realize that their current business could be subsumed by the other when equipped with a new digital infrastructure. With approximately equal cost, each of them can create a network which serves both voice and video as well as new advanced services.

Initially these networks will be moderately expensive (although still fungible). Over time, the absolute cost and the price/performance ratio - i.e. the dollars per megabyte per second will plummet. Physical installation costs and regulatory barriers will be roughly constant (or may change somewhat) but ever higher bandwidths will keep the price/performance ratio dropping even after the absolute cost flattens out. As long as customers have a use for the bandwidth this means that new network will enjoy a big advantage over older ones.

This poses an interesting problem for a carrier - do you install systems aggressively, even though they will be far more costly and less capable than what is just around the bend, or do you wait until they are cheap? If do you wait, a competitor may have staked out a strong position before you do and own the market. If you rush in, you find that a later competitor can undercut you with newer, cheaper systems. In the past, government monopolies have allowed people to have a certain amount of faith in investing in network infrastructure, but that is far from certain.

In the long run the increase in price/performance still poses a dilemma - how can you keep finding services which will utilize the bandwidth? The situation is similar to that discussed above with CPU cycles and storage. The first answer is richer data types, with high quality video being the key example. Notice however that this poses a big problem with respect to existing services - the fundamental cost of a voice call is tiny on a network built for video. The first generation interactive TV networks will carry switched data at 4 megabits/second, or about 1000 times faster than a comparably compressed voice call. The cost of the video must also be lower; about 50 cents an hour at the most in order to allow a movie to be priced competitively with video rental stores. Voice calls which are tariffed (for long distance and business calls) are more like \$6.00 to \$18.00 an hour within the US and far more outside. Future generation systems - both local area and long distance will only widen this gap. An OC-3 ATM network is about 40,000 times faster, and an OC-12 network is 160,000 times faster than voice. As a result, voice costs can't help but trend toward zero as these performance levels become cheaper and cheaper.

Price is not the same as cost of course, but unless regulatory agencies explicitly permit voice to have a different billing structure, all voice calls will be essentially be free. Given any sort of competition at all, existing levels of service will become a lost leader offered to build share for other services. The current telephony world does this with residential voice service in the local area which is given away free in order to create terminals which can be used for long distance. One can easily imagine a situation where all voice services, including long distance, are free if you are willing to sign up to rent a high bandwidth smart video phone.

The long term strategy of every sophisticated communications company that we have talked to is to supplement their fundamental network business with higher level contributions to the overall "food chain". In the case of cable companies the natural move is into programming. TCI, as a case in

point, has used its position and cash to buy its way into minority investments in a number of programming assets - *The Discovery Channel, CNN, Black Entertainment Network, QVC, Carelco* (the movie studio that made *Terminator 2* and *Cliffhanger*) and others. This just covers the current world - they have even more ambitious plans for gaining a piece of the new world created by the information highway. Time Warner is already the world's largest media company and plans to seek synergy with its cable business. The smarter RBOCs have similar ambitions. All of them are interested in using their role in the network business as a point of leverage to build empires which participate in the businesses enabled by the full service network. Their bet is to become information highway companies in the broadest sense of the term, *not* just communications carriers.

It would be misleading to characterize this as a retreat or flight from the communications business itself. These companies will make enormous investments in creating the full service broadband network. This will be a profitable business - at least in the medium term - but a major component of its value will be to put them at the center of the new digital world. Companies that fail to take advantage of the strategic opportunities afforded by this position are likely to be in a very hard spot indeed, caught between the downward spiral of communications cost on one side and diversified highway companies on the other.

## **Retail Banking & Financial Services**

Banking may seem like an unusual industry to pick on, but it illustrates one of the important themes of the information highway. There are 14,000 banks in the United States which cater to retail customers. Their business is based on offering a variety of services - savings accounts, checking accounts, credit cards, loans and so forth.

What is the unique value proposition that these banks offer their customers? The historical answer is geography - most people bank with a firm which has a branch office near their home or commuting path. There are some minor differences in interest rates and other retail banking products which might shift people between two local banks, but hardly anybody is going to find these features sufficiently compelling to make them pick a bank that is 10 miles out of their way as long as there are closer alternatives.

The information highway makes geography obsolete, especially in the case of banking or other financial services which are really just information businesses. The only physical transaction that needs to be done is receiving cash (which itself will one day go away), but automatic teller machines have become the primary way to deal with this. Automatic teller machine networks (note that this use of "ATM" and "ATM network" is quite different from Asynchronous Transfer Mode discussed above) are already threatening the banks. After all, if I have a card issued by a nationwide ATM network, and direct deposit of my paycheck into my account, why do I need a physical bank at all?

In the long term, the notion of depositing money "in" an account is pretty silly. A single transaction clearing service should be able to contract with both me and my employer so that my salary automatically flows out to fixed expenses (mortgage, utilities), or into various short or long term investment accounts. Most purchases would be handled with a card (or digital wallet) which was able to combine the features of a credit card ATM card and digital checkbook.

This will be a huge challenge to existing retail banks. If you take an efficiency perspective is that they should simply close - there is no reason to have human tellers, the "vice president" of the branch and so forth for most transactions. The successful financial service networks will be nationwide in scope, since you can't be otherwise on the highway, and will be highly competitive as

a result. The last vestige of local banking will be the ATMs, which will continue to proliferate. Everything you can do with an ATM except actual cash delivery will also be possible with your PC, your interactive TV and your digital wallet; all equipped with wired or wireless connections to the highway.

Small local banks are not the only ones who will be put to the test. The ability of the highway to host efficient markets will change other aspects as well. Most retail financial services exploit the fact that consumers cannot access the same markets they can, giving them an opportunity to make an arbitrage profit or vig. They also use time access to a consumer's money, known as the "float", to profit from it while the consumer doesn't notice - this is in effect a vig of 100% of the interest they get for a short period of time. The whole notion of savings accounts (which typically have a large interest vig) and checking accounts which don't pay interest (and thus are an opportunity for float) are designed to give consumers a raw deal, and force them to cope with the artificial distinctions and incur overdraft penalties etc.

Money market accounts which issue checks and credit cards show that there is already an effort to cut a better deal for the consumer. The information highway will greatly accelerate this by creating opportunities for far more efficient access to financial markets, eliminating most opportunities for companies to vig the float on their retail customers.

The highway will be the center for financial markets offering of all sorts of new trading vehicles. Even an individual consumer's banking transaction could be traded. A deposit into a savings account is really a case of the consumer making a bid for a specific short term investment with standardized terms. Depositing \$100 into the highway would simply put a bid out to buy \$100 worth of this trading vehicle or security at market prices. One could imagine a variety of firms in a never ending on line auction for your business, so that no two deposits were ever serviced by the same firm. This is already the case with buying and selling stock. As long as the transaction cost is low, why not take this approach on a wide variety of different services?

Services that try to float your money will have to compete with those that rebate a portion of the float income back to the consumer. Aggregation services will exist to bundle small transactions into larger ones which can be directly entered on the open market. Fees will be generated for all of this, but the fundamental fee structure will be based on fairly efficient open competition rather than the closed, inefficient world of today's banking and financial services.

Note that this does not mean that financial services companies will go broke. The basic economics will change dramatically, but the total volume of transactions will go up. The basic opportunity will be very large and some players will make a lot of money in the absolute - although at smaller margins than today. Despite this, many of the current leaders will not survive this transition.

## **Wall Street**

Retail stock market services will fundamentally go the same direction as retail banking. The notion of trading at a single physical location is a historical anachronism that will not survive the transition to the information highway. As various equity and capital markets become decentralized, they will also become more efficient and the various market makers in place today will seem more and more like useless parasites. There is simply no reason for any human stock and bond brokers to be involved in making a market or the process of transferring ownership financial trading.

Once you eliminate the exchange and notion of buying or selling through a broker, most of the other Wall Street jobs that remain will still be in high demand. In fact, the information highway will

renaissance in speculation and investment because it will become so easy to create new trading vehicles and markets. Mutual funds and other aggregated securities will sprout up overnight and huge new markets will be developed.

Analysts will find a zillion new markets and trading vehicles to study and pontificate upon. It would be nice if I could predict that market analysts would actually know what they are doing, and be subject to proof of competence, but alas I doubt this will be the case. Unless there is a breakthrough in economics, market analysts are likely to remain at a level just beneath chiropractors and new age faith healers in terms of their efficacy.

## **Print Based Media**

One of the groups most concerned about becoming roadkill are the people involved in wide spread analog information distribution - newspapers, magazines and books. Each of these will suffer a somewhat different fate.

Books are actually the best off of the bunch. I believe that print based books will continue to be published for the next 50 years, and possibly longer although by the end of this period the volume will be far smaller than today. Books have a number of things in their favor versus the information highway alternatives which other forms of print media do not have.

Customers *pay* for books. This may sound like a silly distinction, but in most other print media *advertisers* pay for a great deal of the cost. Advertising revenue is far more fickle than direct consumer demand, and information highway has many features which will be very tempting for advertisers. Inroads that the highway makes in these areas will hurt other media, as discussed below, but this will leave books unscathed.

Another advantage to a book is that they are long and, with a couple of exceptions, they are not suitable for random access. People tend to read them all the way through, rather than browsing selectively, or starting in the middle. This has two consequences for competition with the highway. The first is that the quality of the display is very important. You might put up with a computer display to check stock prices, or see the latest news, but choosing to read the next Tom Clancy thriller on line will mean signing up to reading several hundred pages on the computer, PDA or television. Current display technology is *not* good enough for people to do this because of limitations in resolution, contrast, weight and power consumption. Ultimately this will be solved, but it will give books five to ten years longer than other more casually consumed print material will have.

The second advantage to length and uniformity is that most books do not benefit as much from indexing, retrieval and filtering. On line newspapers or magazine articles can be browsed, filtered and indexed on a computer in ways which are far superior to print. This is also true of some kinds of books, particularly reference books such as dictionaries or encyclopedias which will rapidly move to the highway. Novels and most non fiction books do not benefit as much from this technology. It might help you find the book, but once you have found it you will read for a long time before searching again - the ratio of retrieval to reading is low.

Newspapers are in probably the worst situation of any form of print media. I expect newspapers to be published for very long time, they will *not* disappear overnight. They will face some huge challenges however. In fact, many newspapers are already experiencing problems because of the long term effects of competition from television.

The biggest problem with newspapers is their reliance on advertising revenue. The current budget for newspaper advertising is (in adjusted dollars) half what it was in 1950. The obvious reason is that

during the same time television has become a far more important venue for national advertisers. This trend will continue because the highway will continue to suck advertising dollars at both the high end and low end of the market - more effective and extensive coverage for national advertisers, and at the same time more cost effective ads for small business. Newspaper ads from businesses tend to be from local companies, at a scale which cannot be addressed by television. This will change as the information highway allows local ads to be cheaply created on advanced personal computers and delivered inexpensively on the highway. As local firms shift their advertising dollars to the highway newspapers will feel the pinch.

The long term decline in business advertisement has already shifted the burden to classified ads, because they were immune to competition from television. Classifieds are the largest single source of income to most newspapers, comprising between 50% to 80% of revenue. Unfortunately for them, the information highway will have *many* alternative ways for people to sell small volumes of merchandise, find jobs and other classified ad functions. In some cases, such as used cars, the highway gives a mechanism for people to create an on line auction - not only attracting customers, but also offering better pricing. More generally, classified ads are inherently a database problem - they aren't about great graphic design, or catching the eye of a casual reader; in fact they are *not* very user friendly. They exist because enough people are sufficiently motivated to laboriously wade through the fine print in search of the car, home or job that they want. People *that* interested would be a lot better off registering a filter query or browsing the highway from their PC, interactive TV, smart phone, or perhaps an interactive kiosk.

A decline in classified revenue and accelerated loss of business advertisement would put newspapers under large financial burdens, even if the core editorial side of business did not have much direct competition from the highway. The inevitable result will be higher prices for the customer - up to *five times higher* if they are going to foot the bill entirely. This kind of increase will hurt circulation. Fixed costs, like maintaining the editorial staff and foreign bureaus will have to be born by a smaller clientele, increasing prices further. The big problem with newspapers isn't the news itself - it the way it is paid for.

This isn't to say that there will be no competition for news however. There will be a lot of it on the highway. Today's major newspapers, like the *New York Times* or *Washington Post* maintain international bureaus and put a large premium on having better in depth coverage than can fit in a 30 second *CNN* spot, and they staff accordingly - an editorial staff of about 1000 for the *Times* alone. The forest of microphones and video cameras thrust in the face of newsmakers around the world is a testimony to incredible redundancy in the way news is currently gathered - major newspapers, weekly news magazines, TV and radio news all have created redundant news gathering organizations. As the information highway coalesces the various media, this redundancy is likely to decrease substantially.

There will always be a market for both depth and breadth, and the information highway provides an ideal way to distribute news. There will be a number of ways this will happen. One natural direction is to make digital editions of existing newspapers. Initially this will be done with repurposed content taken from the print edition, but in order to really take advantage of the new properties of the highway, subtle but important differences will be necessary. Newspapers must cater to a very broad set of customers, but on the highway that is not necessary. Each user could in principle get a different edition, by turning their specific interests into criteria for selecting articles - both by posting filter queries and by giving the human editorial staff more input. The queries can include high level interests (skip the sports page for me), down to more specific items such as news on your company, or home town.

Text based news has the advantage that it can be sent to existing PCs, AtWork FAX machines and smart phones and over narrowband wireless links to PDAs. The disadvantage is that the more text oriented a service is, the less well adapted it will be for interactive television. This brings up the other natural direction for news is to do a similarly extend the depth of video news - breaking it into segments which can be queried, filtered and assembled into custom editions in a similar manner to text based newspaper articles. Text and graphics can be also be used as annotations to add depth to what is fundamentally a video based service. It will be interesting to see what mix of text and video, and what mix of existing text and video companies, will eventually win out as the primary form of news on the information highway.

The fate of magazines is somewhere between that of books and newspapers. The business side is intermediate - they typically get half of their income from advertisers, so losing them would not be a lot of fun, but the nature of the advertisements is usually tied to the special editorial focus of the magazine and thus is less vulnerable.

The editorial side is more interesting. Magazines have a key characteristics which will also be found in special interest group publications on the information highway - tight editorial focus on a particular topic. A subscriber to *Skin Diver*, *Rolling Stone*, *Communications of the ACM* or *Food Arts* has, in effect, already posted a filter query on a range of interesting topics. Common topics like "cars" have already been subdivided into smaller domains - *Road and Track*, *Hot Rod*, *4WD*, *Hemmings Motor News* and at least a dozen other magazines have fine tuned natural boundaries in people's interest into their own editorial niches.

As a result, many on line publications on the information highway will be modeled on magazines to one degree or another. This will not spell instant doom for print magazines, because many of them will be able to leverage their brand name and customer base over to the new medium. Print and on line versions could coexist for a very long time and it wouldn't be surprising if the word "magazine" was coopted by the information highway to apply to the new format and a couple decades hence people will have to explain to their wide eyed grandchildren that, once upon a time, magazines were delivered once a month on *paper*.

## **Broadcast Television Networks**

Television broadcasters are some of the most likely fodder for roadkill of any of the current media companies. Some form of broadcast TV will continue on for the next 20 to 30 years, but its pivotal role as the leading form of mass media is already declining due to competition from cable and other factors. Their decline will be long and slow and may take the next two decades, so this isn't an overnight phenomenon, but by the time that they realize it the window of opportunity for changing anything will have long since passed. The role of TV networks in the world will be tremendously effected by the information highway, yet they are doing very little about it - like a dazed rabbit caught in the headlights of an oncoming car.

As in other areas discussed in this memo, the trend is already well under way. Cable TV relived the pressure on channel space, and thus enabled special purpose television. Prior to this the opportunity cost of time on one of the three major broadcast networks was so high that the networks were forced to cater to the same set of customers. It made more economic sense to pander to the lowest common denominator and pull in the most viewers, even if that wound up making the networks quite similar and the shows quite moronic. Each had a similar line of up shows and none of the stood for a particular set of viewers or values. Each of them wound up having specific areas where they would happen to create a long term asset which would run for

twenty years or more - NBC with the *Tonight Show* and *Saturday Night Live*, CBS with *60 Minutes* and ABC to a lesser extent with leadership in *ABC Sports*. Apart from those few long term successes which were basically unique to them, they had a very comparable crop of cartoons on Saturday morning, soap operas and old movies during the day, network news at around 7PM, sitcoms and dramas in prime time and so forth. These programs would only last for a few years before going into syndication because they were created and owned by third parties (via the recently repealed "Fin-Syn" regulations).

This situation was possible because they had a triopoly with very limited resources. There are only 21 hours of prime time per week and only three sources of new programs at any one point. Nobody else commanded enough market share to host new programming on a broad scale. Independent TV stations wound up showing syndicated reruns that was about it. Although the individual companies would have their ups and downs (witness NBC in recent times) being in a three way race meant that you would always win, place or show, and thus would always be in the money to at least some degree.

Critics might complain that there was never anything on TV that they wanted to see, but the job of the networks was to try ensure that, statistically speaking, this proposition was false for the largest possible number of customers. A major newspaper must do the same thing - attract lots of customers under the constraint that they have only one product and only so many pages to go around. This is not the way that the book business works of course, because there is no artificial constraint on the number of publishers, or the number of books equivalent to the three networks and 21 hours of prime time. There are very few barriers to entry in publishing when compared to network TV. Many books are written and a number become successful, but those are not the same ones that would have been picked up front to succeed.

My favorite example of this is *A Brief History of Time* by Stephen Hawking. It has sold 5.5 million copies worldwide and was on the *New York Times* Bestseller List for months. If publishers had to pick their "fall line up" the way a TV network does, and made a huge opportunity cost bet on selecting one book versus another, who would have picked a physics book? This isn't an isolated example, *The Bridges of Madison County* is the first novel by an English professor. It received essentially no advertising budget and was not expected to sell beyond the initial run. Word of mouth, particularly through book stores, put it at the head of the bestseller list where it has stayed for ages. Meanwhile many highly vaunted and hyped books fail.

Cable broke this triopoly by allowing many more channels. This pushed the economics past a critical point where channels were forced to emulate each other; instead they had to differentiate and have a unifying theme. Specialized networks like *CNN*, *ESPN*, *MTV*, *Nickelodeon*, *The Discovery Channel* and others have fractured the uniformity of the TV audience. The information highway will take this step much further because it will break the constraint of a fixed schedule - any show or content can be watched at any time. This furthers the fragmentation of viewership to the point that it looks very much like books or magazines - thousands of potential choices, with little or no opportunity cost in making them available.

In addition, the highway will break the traditional "follow on" effect that comes with fixed channels. It is a well established principle in the TV industry that if people watch one show they are likely to just keep watching into the next one on the same channel. The people who create the schedules try to line up shows to encourage this and "build a strong Thursday night" and use a few hits to prop up a mediocre set of offerings. Video on demand breaks this cycle - you can watch anything you like at any time. Personal agents might recommend shows to you in any sequence. The whole notion of a channel simply disappears.

Another troubling issue with the information highway is that it may suck advertising revenue away from the networks. Advertisers on broadcast TV (networks plus independents, but not cable) currently pay the equivalent of \$21 per month per TV household. That is really a lot of money considering that TV is such an indirect means to achieve a sale. Ads on the information highway will merge with shopping services in that they will give the viewer the ability to order directly, switch to an in depth infomercial or other avenues which are not available in standard broadcast TV. In the next five years this will not constitute a real threat to the networks as a major market for advertising, but in the long run it will have a significant effect.

The TV networks do *not* know how to respond to this threat. I have discussed the trend at length with senior network executives and independent TV producers. Some see the threat, and others shrug it off. The fragmentation of the TV audience is viewed as a major disaster, akin to an epidemic. Some claim that the public won't accept changes from the established model of TV, despite the rather clear evidence from cable TV, video rental, books and magazines that once given a choice people tend to take you up on it. Others admit that it is happening, but don't know what to do about it. As the sit wide eyed and dazed, the wheels are almost upon them.

## **Hollywood**

The motion picture industry has developed a very predictable reaction to technology. In the past 50 years there have been a number of technological innovations - talkies, color films, television, VCRs and cable TV. In each case the initial reaction was the same - the new breakthrough was the work of the devil, an evil force which would destroy the industry. These fears have never proved out, and in fact just the opposite has happened. Each change caused a large expansion of the market for feature films and greatly increased revenues.

This rosy picture may well happen once again. Movies are likely to be popular for at least the next 20 years, and probably far longer. The information highway is a great way to deliver films - in fact a far better one than the current schemes. Hollywood's fortunes will continue to improve as the highway takes a larger share of film distribution.

Movie theater attendance is continuing to drop, to the point that theater revenue is no longer the largest source of revenue for most films. The good thing about theaters for Hollywood is that the price is relatively high, and the margin is good - typically about 50%. Despite this, the current trend is likely to continue and theaters will become a smaller and smaller phenomenon, although they are unlikely to become completely extinct. In fact stage plays are a good model - they are still around despite movies, television and other art forms. Plays no longer hold the same relative position as a medium of mass entertainment, and one day the same thing will be true of movie theaters. They will be expensive places where people purists, aficionados and those with nostalgia go for a special night out.

Blockbuster and other video rental stores are currently the largest source of revenue in absolute terms. Their business is almost certain to go away as the information highway provides large libraries of films on line. However unfortunate this is for video store owners, it is a boon to Hollywood because the current video rental business is based on selling the tapes to the stores rather than directly participating in the rental income. Video on demand on the highway will be done as a share of revenue which is far more lucrative for the film's owners. The death of the video store will greatly increase the bottom line to Hollywood.

This covers existing movies, but what about the new media? The role of Hollywood in the new

business is much more uncertain than the distribution question. It is far from clear what role the major studios will play, if any, since in any direct sense they don't really make the movies - they finance them and arrange distribution deals and perhaps rent them some production facilities. The actual creative work is done by a collection of individuals - actors, directors and employees at small companies like Industrial Light and Magic, which come together for an individual project. Financing and distribution will still have to occur, but they are as likely to be done by integrated highway companies as they are by decedents of today's studios.

As the information highway and the growth in computing price/performance progress, new narrative and entertainment formats will develop. Computer games will increase in production values until you won't be able to tell the difference between the game and a movie - they will be equally realistic. High bandwidth communications over the highway will enable multiple people to come together and share the same experience - real or simulated. Who will create the new media formats and entertainment for the information highway? Many of the creative people will come from the community that currently makes movies, but I suspect that the majority of current Hollywood talent will not make the transition.

The most dramatic shift for Hollywood will be the change in production costs and equipment. Over the course of the last 10 years publishing has moved from having an incredibly expensive analog production process - up to thousands of dollars per page - to one where a PC and software are sufficient for both professionals and amateurs alike. The same thing will happen to video editing, post production and special effects, as the growth in computing capabilities continues. The amazing dinosaurs in *Jurassic Park* will cost 1000 times less in ten years, making this sort of effect available to nearly anyone. A few years after that and effects of this capability will be in children's toys - the ultimate extrapolation of the Etch-a-Sketch.

## **Personal Computers**

I've saved the best for last. Our own industry is also doomed, and will be one of the more significant carcasses by the side of the information highway. The basic tasks that PCs are used for today will continue for a long as it makes sense to predict, so it isn't a question of the category disappearing. The question is one of *who* will continue to satisfy these needs and how?

As a case in point, consider that the fundamental category needs for mainframes and minicomputers *also still exists* and will continue to do so for a very long time. Despite this, the companies involved are dying and the entire genre is likely to disappear. The reason is that a new breed of machine - the PC - came along which out flanked them. In the early years PCs were not particularly good at what minis and mainframes did, but they were terrific at a whole new set of problems that the traditional computing infrastructure had basically ignored.

Personal productivity applications drove PCs onto millions of desks and created a very vital industry which grew faster - both in business terms and price/performance - than the mainframe and minicomputer markets. The power conferred by this growth made PCs the tail which wagged the dog; free to ignore the standards which existed for mainframes and minis and move off on their own. Over time the exponential growth in computing has finally (after 17 years) given the PC industry the technical ability to beat minis and mainframes in their own domain. Although the early software platforms for PCs had to be extended to fully realize this potential (Dos to Windows to NT to Cairo), it turned out to be far easier to do this than to make mainframe or minicomputer systems address the new needs and applications. Even within the heart of minicomputer and mainframe's domain - giant transaction processing applications etc., the old standards will not be used.

I believe that the same thing will happen again with PCs playing the role of mainframes and minis, and the computing platforms of the information highway taking over the role of the challenger.

The technical needs of computers on the information highway, or IHCs are quite different than for PCs. The killer applications for IHCs in the early years will include video on demand, games, video telephony and other distributed computing tasks on the highway. It is hard to classify this as either higher tech or lower tech than the software for PCs, because the two are quite different. Most IHCs will certainly need to be cheaper than PCs by an order of magnitude and this will inevitably cause them to be less capable in many ways, but some of their requirements are far more advanced.

Another way to say this is that the rich environment of software for PCs is largely *irrelevant* for IHCs. Windows, NT, System 7 and Cairo do *not* solve the really important technical problems required for IHC applications, and it is equally likely that the early generations of IHC software won't be great platforms for PC style apps. This isn't surprising because they are driven by an orthogonal set of requirements.

The IHC world will almost certainly grow faster than PCs, both in business terms and in price/performance. The PC industry is already reaching saturation from a business perspective. Technically speaking, the industry is mired in hardware standards (Intel and Motorola CISC processors) with growth rates that are flattening out relative to the state of the art - just as the 360/3090 and VAX architectures did. The Macintosh and Windows computing environments may be able to survive the painful transition to new RISC architectures, but they will lose time and momentum in doing so.

PCs will remain paramount within their domain for many years (we'll still have a computer on every desk) but IHCs will start to penetrate a larger and larger customer base on the strength of its new and unique applications. The power of having the worlds information - and people - on line at any time is too compelling to resist. For a long time people will still have a traditional PC to handle traditional PC tasks - in precisely the same way that they have kept their mainframes and minis for the last 17 years. One day however people will realize that their little IHCs are more powerful and cheaper than PCs - just as we have finally done with mainframes. There will be a challenge for the IHC software folks to write the new systems and applications software necessary to obviate PCs, just as we had to work pretty hard to come up with NT, but this battle will clearly go to the companies who own the software standards on IHCs. The PC world won't have any more say about how this is done than the companies who created MVS or VMS did about our world. Of course, some of the VMS *people* were involved, but as discussed above it is very hard for *organizations* to make the transition.

This may sound like a rather dire prediction, but I think that for the most part it is inevitable. The challenge for Microsoft is to be sufficiently involved with the software for the IHC world that we can be a strong player in that market. If we do this then we will be able to exploit a certain degree of synergy between IHCs and PCs - there are some natural areas where there is benefit in having the two in sync. The point made above is that those benefits are not sufficiently strong that they alone will give us a position in the new world. We'll live or die on the strength of the technology and role that we carve out for ourselves in the brave new world of the information highway.

-----

<sup>1</sup>Obviously, these benchmarks are not fully equivalent and "your mileage may vary", but the trend is still correct within a factor of two one way or the other.

<sup>2</sup>Depending on whether the universe is open or closed. These are the same alternatives that T.S. Elliot presented - either a bang or a whimper.

<sup>3</sup>This may seem self serving, but its a good time to point out that I used to be a cosmologist myself!

<sup>4</sup>After all, Dom Perignon invented champagne, and colleagues of his gave the world brandy, Roquefort and the science of genetics!

<sup>5</sup>It's not just a good idea, it the *law*!

<sup>6</sup>The science fiction novel *Ender's Game* by Orson Scott Card explores this notion.

<sup>7</sup>Well, there goes any hope of *me* every getting a cabinet seat!

-----

Microsoft Confidential