

# Security and the Information Highway

Craig Dobson, *AGT Ltd.*<sup>†</sup>

**Abstract**— The developing information highways will fundamentally change our world -- but only if they can be kept secure. As the fundamental premises upon which these highways are being built run contrary to those needed to ensure its security, this paper provides some insight into both these premises and their implications with respect to how these highways are likely to evolve.

## I. INTRODUCTION

**S**ECURITY has a problem -- the fundamental premises upon which the developing information highways are being built run contrary to those needed to ensure its security. Fast PCs, open systems, interconnectedness, distributed processing, virtual circuits, convergence, portability, global reach, advanced user interfaces, agent-based technologies, and so on, enormously complicate the task of ensuring the integrity, safety, and security of both the infrastructure and its 'clientele'. As the developing infrastructure becomes more pervasive, the application domain, i.e., the activities that it will support and enhance, is increasing significantly. Some 37 million computers already access the Internet on this continent alone, the information available on the net appears to be growing exponentially, and electronic fund transfers exceed a trillion dollars a day. The more we become entangled by these electronic webs, the more ways we can be hurt by them, and the more issues that need to be resolved. While issues such as virus attacks and privacy make the news regularly, others such as intellectual property protection (particularly once electronic cash becomes commonplace) are equally difficult. As well, the age-old problems indigenous to the security arena have not gone away. Adding security is a dynamic and still largely reactive endeavour that is more of an insurance policy than a revenue generator. Security is often a matter of intent and the costs must always be balanced against those needed to bribe an appropriate employee or those of restoring a system from backups.

The remainder of this paper deals largely with the first issue. By providing you with a sketch of what is to come in the telecommunications and information services arena, you will hopefully be better able to appreciate what you are in for. The opportunities and potential of a true information highway are enormous. They will only be truly realized, however, if it can be kept secure.

<sup>†</sup>Mail: 20E, 10020-100 St., Edmonton, Alberta T5J 0N5  
Phone: 403-493-3050; Fax: 403-493-4277  
Internet: cdobson1@rnd.agt.ab.ca

## II. ENABLING TECHNOLOGIES

Making sense of the changing telecommunications and information services arena can most easily be done through an examination of the implications of the relative rates of progress being made in the underlying enabling technologies. The current rate of progress in five key technology areas over a ten year period is depicted in Fig. 1. With integrated circuit densities doubling every eighteen months and the processing power associated with the resulting microprocessor chips increasing even faster, it is becoming economically feasible to convert all signals to digital form. With this digitization comes media independence, multimedia, and the associated notion of convergence. Convergence will affect content, terminals, and networks, and significantly blur the boundaries between computer and telecommunication systems.




Figure 1. Enabling technologies.<sup>1</sup>

Another aspect of more economical and powerful processing power is the some thousand-fold increase in spectrum capacity that becomes possible. With such improvements, wireless network equivalents for telephony, data, and cable are becoming available -- simultaneously increasing both the number of network access points and the network's intrinsic complexity.

The photonics group of enabling technologies underlying the four hundred times improvement in lightwave capacity includes fibre optics, erbium doped fibre amplifiers, and soliton transmission. Due to these developments, the cost structure of the telecommunications industry is being turned upside down -- whereas long distance service used to be expensive, and local service cheap, the situation has now reversed itself; whereas switching costs have decreased only slightly in the past 25 years, long distance transmission costs have been declining exponentially. Data transfers and remote login via the Internet are now largely usage and distance insensitive,

and any network connected hacker essentially has the world at their doorstep.

Progressing at a somewhat slower pace are the user interface technologies such as speech recognition. Finally, after more than a decade's worth of intensive effort, the artificial intelligence (AI) -based engines needed to bring these technologies to fruition -- and simultaneously make computer power accessible to the masses and forever pacify concerns about computer illiteracy -- are finally beginning to mature. While the agent-based technologies that will ensue will greatly enhance and increase our interaction with computers, they will also open the door to a myriad of ills from a security perspective.

With hardware technology now relatively cheap and pervasive, the key to future success will lie in our ability to develop and utilize software. Unlike the success stories associated with the hardware-based technologies mentioned above, industry's ability to develop software is embarrassingly slow and it is largely in an effort to circumvent these problems that systems were distributed -- the mainframe to PC phenomenon that spawned the software interoperability problems that spanned the decade. New approaches and languages now portend a solution. As these take root, our interconnectedness will increase -- Internet is already growing at a rate in excess of 100%/yr -- as will the security headaches.

### III. THE INFORMATION HIGHWAY IS HERE

Victor Schnee has pointed out that the *'information highway' is more akin to the opening of a public beach on the ocean of information than it is to any highway currently in existence.*<sup>2</sup> Nonetheless, the information highway is the accepted term, and is the term that is used in what follows.

*For the past few years the titans of media and communications have waged a war for the digital future. // Shambling towards their distant goal of a wired world, they have been too busy to notice the unruly bunch of computer hackers, engineers and students scurrying about at their feet. They should have paid more attention. For while the giants have just been talking about an information superhighway, the ants have actually been building one: the Internet. // Last year the Internet as a whole doubled in size, as it has done every year since 1988 [see Fig. 2]. At this rate, within two years the citizens of cyberspace will outnumber all but the largest nations.*<sup>3</sup>

Prior to discussing how the information highway infrastructure might evolve, let us look briefly at what the highway might carry. Bypassing the obvious 'multimedia' answer, one wonders about voice and video. After all, the Internet is world-wide and multimedia does include audio and video, right? Boy, avoiding long distance charges would be nice!

Regardless of the underlying infrastructure used to support the information highway, there is a significant difference between the voice and video traffic carried on current telecommunication and cable television networks and the data traffic carried on the Internet. The difference has to do with

the time sensitivity of the data streams. Whereas Internet is a store-and-forward network, voice and video networks operate in real-time. Store-and-forward networks sequentially forward and store the data in multiple computers enroute from the source of the data to its destination. Intermediary computers only forward received transmissions when they have time. Receiving pieces of a phone call, a few bits at a time and with varying delays between the bits, does not do much for a caller's patience. In real-time networks, the data streams are continuous and this does not happen.




Figure 2. Growth in host computers on the Internet.

While transmitting voice and video over the Internet is currently possible, the quality is poor and the required computer equipment at each end rather prohibitive. Bandwidth on the Internet is increasing, though, computer costs are coming down, and work is underway to provide more real-time services.

### IV. EXISTING INFRASTRUCTURE

Given that the Internet will form the basis on which the information highway will grow and prosper, one wonders how the currently bandwidth limited Internet will evolve. Though the infrastructure associated with a 'converged' multimedia network or information highway is far too expensive and would take far too long to build from scratch, the required infrastructure could be synthesized by utilizing the appropriate components of each of the three existing groups of players -- the telecommunication, cable television, and information services industries. Built to optimally serve the requirements of their primary applications suite, the three networks deployed are both different and complimentary. At a high level, the three networks are depicted in Fig. 3.

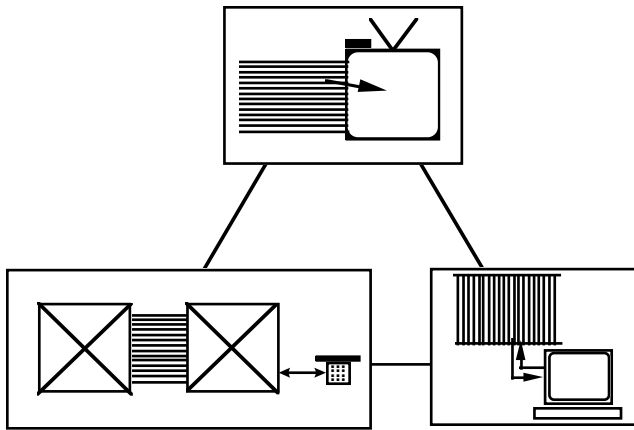


Figure 3. Existing networks.<sup>4</sup>

Current telecommunication architectures are analogous to those of mainframe computers in that all intelligence is highly centralized and used to support a myriad of dumb terminals (in this case, the telephone). The boxes appearing in Fig. 3 represent the specialized supercomputers (switching machines) needed to route and bill calls appropriately. Whereas the facilities interconnecting these systems are mostly digital and run up into the Gb/s range, the links running to the home are mostly bi-directional 3.3 kHz analogue links. Due to the narrowband nature of the latter, only that information (or phone call) intended for a particular telephone travels to that telephone. To concurrently reduce costs and increase flexibility, the telecommunication's industry is off-loading their centralized intelligence in two steps. First, the switch software is being split in two -- that which is service related and that which is strictly switch related. While these are currently referred to as the *advanced intelligent network* (AIN) developments, the longer term approach falls under what is known as *telecommunication management networks* (TMN). Subsequently, with user to user signalling becoming available, many services -- such as call forwarding for example -- will be off-loaded to users' terminal equipment. Though this will greatly facilitate work in the *computer telephony integration* (CTI) field, the security implications of allowing third party access to telco signalling networks are obviously significant.

CATV networks, on the other hand, currently have little intrinsic intelligence in either the network or the end terminal (the television): they currently consist of a very wideband unidirectional analogue facility running from a headend to every home in a serving area. Since all information (TV channels) required by all connected homes travels on this distribution facility to all connected homes, it is up to the converter box on individual television sets to select and display only that channel its viewer wishes to watch.

In LAN environments, intrinsically dumb bi-directional wideband digital facilities link intelligent terminals. All information to and from all terminals travel on these shared facilities and it is up to the end devices to sort out which messages are for them and which are not, as well as to know how format and address the messages they wish to send.

Bridges, routers, and gateways have been developed to interconnect these networks and it is largely on this framework that the Internet has been built. That these architectures do not scale well relative to switch-based ones has not been too apparent due to the currently 'thin' or low capacity narrow-band store and forward traffic these networks typically carry. This is changing, however, as the user base expands rapidly, the traffic becomes more multimedia in nature, and the real-time requirements are increased. To counter these issues and help stabilize the network, operators are beginning to centralize both processing (e.g., client server architectures) and routing (virtual LANs) activities.

Relative to the developments outlined below, these networks represent the 'good old days' from a security perspective. Centralized digital telecom networks with out-of-band (inaccessible) signalling and dumb telephone sets, like the mainframe environments of years gone by, are pretty ideal -- apart from some long distance theft, they are almost impenetrable. In coaxial cable television distribution systems, there is no intelligence and nothing to steal. Hence, the problems that you have most likely been concerned with are those associated with data networking. While those represented by standalone LANs were straightforward, now that they are being interconnected on a grand scale, things are starting to change -- but that is part of the discussion to follow.

## V. WIRED OPTIONS

A fully capable information highway might be assembled from the access infrastructure of the cable industry in the residential areas, LAN access in the business areas, and the interoffice, interexchange facilities of the telephone companies. To make this plan successful, the telco switching machines would have to be bypassed to alleviate congestion (allow unlimited connection times) and provide additional bandwidth, the CATV infrastructure would need to be made bi-directional, and the LAN technology would need to be modified in order to allow it to scale.

Since this solution is unlikely, due in large part to regulatory constraints, there are a number of developments underway to permit each network player to provide broadband data services independently by augmenting their own infrastructure -- see Fig. 4. With respect to access, LANs are okay in that they already exist on high bandwidth infrastructure. Both cable companies and telcos, though, have work to do.

Over the past year, cable companies have been announcing the commercial deployment of cable modem technology. Using upgraded cable infrastructure, these devices will be capable of providing bi-directional data services in the 0.5-27 Mb/s range.

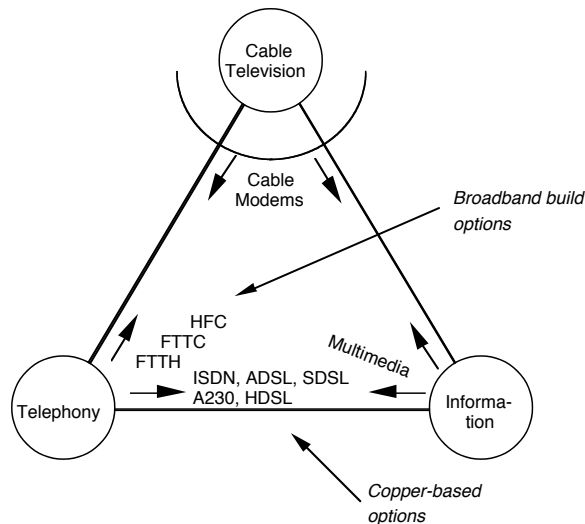


Figure 4. The expanded playing field: FTTC: fibre to the curb; ISDN: microlink services; ADSL: asymmetric digital subscriber line; SDSL: symmetric DSL; A230: Access 230 from TRLabs; HDSL: high speed DSL.

The telephony industry is pursuing two sets of options -- coax/fibre deployment scenarios largely aimed at the full service capability (including broadcast) and the copper-based solutions more commonly focused on data services. ISDN is switched-based (usage sensitive), low bandwidth, and expensive. ADSL, SDSL, and A230 are all data above voice solutions; i.e., coexist with telephony services on the same copper pair. HDSL requires 2-pr copper, which, at least with respect to aerial plant, is not available.

Unlike either the switched ISDN solution or the ADSL, SDSL, A230, or HDSL options in which each user has a physically distinct and separate access line, the cable modem solution shares access bandwidth in the same manner as a LAN. Hence, while a user is only supposed to decode the packets intended for him, all users in the same cable serving area would receive all the data packets -- a less than ideal situation from a security perspective.

Regardless of the technology employed, these access networks must be interconnected via broadband backbone networks -- networks comprised of transport and routing components. While the transport requirements will be handled by fibre-based *synchronous optical network* (SONET) equipment, the evolution of the routing/ switching components is controversial. Though current Internet traffic is largely routed, more scaleable cell-switched *asynchronous transfer mode* (ATM) machines will likely be needed if these networks are to scale sufficiently to meet projected traffic loads. At this point, though, the routing option is more cost-effective and widely deployed while ATM is neither. Matters are complicated by the fact that Internet traffic does not yet happily coexist with ATM-based equipment.

To allow various media types to coexist on the same network, the associated data is chopped up and placed into packets or cells which are then routed or switched through the

network as needed. The more real-time the requirements, the smaller the cells have to be. In the case of the 'grand' compromise, i.e., ATM, the cells are 53 bytes long. While more connection-oriented protocols (virtual circuits) are used to satisfy real-time requirements, connectionless arrangements prevail in more data oriented store and forward architectures. Since in neither case does a physical circuit exist and, in connectionless arrangements, each cell in a particular data transfer might traverse a different route, 'call trace' options are not too practical. These networks have also been optimized for bursty traffic and data transfers and circuits only exist for very short periods of time. Indeed, every transfer from a web site is effectively a different connection -- when surfing the web, typical users hit several countries a minute and potentially hundreds of computers a session. Since these services are provided on a largely usage and distance insensitive basis, essentially no records of these transactions are kept.

In the real world with which you deal, all of these networks will be interconnected and have to interoperate. Indeed, the absurd regulatory environment that is mandating such things as unbundling, interconnection, and resale will force not only interconnection between networks, but interconnection between various components of various networks, including the wireless ones outlined in Sec. VI. Innovative companies will effectively be able to piece 'new' networks together simply by mixing and matching appropriate components of others' networks as they see fit. We will have to make it possible, you will have to keep it secure.

## VI. WIRELESS OPTIONS

With the exponentially decreasing costs of electronic circuitry, it is becoming possible to establish either complete or partial wireless network overlays to either bypass or enhance the wired infrastructures outlined above. This new and expanded playing field is shown in Fig. 5. Mobile, cellular, VSAT, DBS, and TV broadcast systems are well established and will not be discussed.

Originally intended primarily for video signal distribution, the CellularVision wireless cable system introduced some five years ago operated in the 27.5-29.5 GHz band and delivered 1 GHz of bandwidth to every home. In retrospect, the CellularVision system was the first of what has now been termed *local multipoint communication systems* (LMCS). The second generation systems now becoming available from a variety of manufacturers support bi-directional high bandwidth voice, data, and video services. These systems do not yet support user mobility.

When you blow up mainframes, you get PCs and when you blow up big geosynchronous satellites -- *low earth orbit satellites* (LEOS) result. Intended primarily for high speed data transfer to the boonies, world-wide, the Gates-McCaw Teledesic initiative would place some 800 satellites in low earth orbit. The low orbit reduces latency and the large number of satellites provides both the coverage and capacity

needed. It will be at least five years before such systems are deployed.

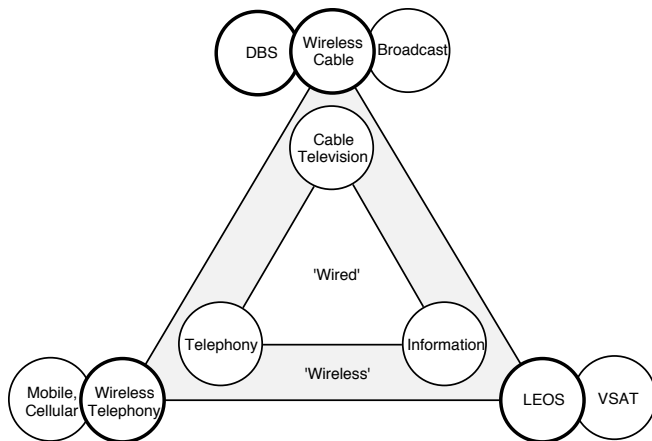


Figure 5. The expanded playing field:<sup>4</sup> DBS: direct broadcast satellite; LEOS: low earth orbit satellite; VSAT: very small aperture terminals.

Wireless telephony refers to *personal communication systems* (PCS) -- small low mobility voice terminals which amount to a cross between a cordless telephone and a cellular handset. They will compete directly with local telephony and will be initially deployed in large metropolitan areas. If these systems do not materialize quickly, they may lose out to advanced digital cellular systems. Though wireless connections and portability create significant problems from a security perspective -- both in terms of access control and the intrinsic complexity of the supporting infrastructure, the management systems associated with these terminals, will be able locate the users to within a few blocks (which raises some privacy concerns).

#### VII. DISTANCE BECOMES IRRELEVANT

Though currently held back by the capability of the associated opto-electronic equipment, the capacity of a single mode fibre optic cable is effectively unlimited (~75 000 GHz). Commercially available systems currently run at 2.5 Gb/s and lab-based systems are some 40 000 times more capable. As these developments are deployed, the incremental costs of higher bandwidth long distance traffic will, as shown in Fig. 6, approach zero. If long distance tariffs become marginalized, many of the value-added applications could be provided through off-shore facilities. On the networking side, this implies fewer highly dispersed large switching machines and more direct call routing. On the usage side, this opens the door to third world 'white collar' telecommuters, ..., and hackers.




Figure 6. The decreasing cost of long distance transmission.<sup>5</sup>

#### VIII. COMPUTER LITERACY IS IRRELEVANT

It is interesting to note that multimedia is a largely unidirectional phenomenon, i.e., from the computer to you. Getting from you to the computer is something less than multimedia, say a mouse click or maybe even a keystroke or two. After many disappointments and false starts, the field of artificial intelligence is starting to mature -- making the way for natural language interfaces and machine learning. The former will facilitate richer human-computer interactions and significantly stimulate the use of computer systems. More importantly, making computer literacy irrelevant will make these highways, together with all their promise, available to a much broader cross-section of society.

When incorporated into knowledge-based systems, machine learning will eventually yield systems with the ability to recognize and analyze their errors so that they can modify themselves to avoid making similar errors in the future. Smaller scale, nearer term uses include the development of personal agents -- sophisticated software programs that can be sent into the network to work on behalf of the user. Example applications include the retrieval and filtering of information, routing, scheduling, and task automation, decision making and planning, and pattern recognition, simulation, and modelling. With negative intent, such capabilities could be used to do significant damage. Instead of remotely doing the dirty work personally, a hacker, for example, might send in an agent in to do it locally and then report back once successful. Many first generation agents are optimized for information gathering and filtering -- the only difference between gathering information on nuclear materials for a high school physics report versus that for making a bomb is intent.

## IX. THE NETWORK IS THE COMPUTER<sup>6</sup>

While Internet may provide the networking infrastructure, and advancing electronic and fibre optic technologies cost effective terminals and distance insensitive networking capabilities, the software interoperability/compatibility problem remains -- and the dearth of improvement in the software development technologies is not helping. The Internet explosion in the early nineties was partially due to a program put out by Netscape (formerly Mosaic) Communications. Other than providing a greatly simplified user interface, the key to Netscape was that, for the first time, it laid out -- from the top or network side down -- a *hypertext markup language* (html). Anyone wanting to make documents available at a web site, regardless of their home terminal's capabilities or operating system, simply had to use html to define their document. Conversely, anyone wishing to view or print the document, simply needed to have an html browser available on their home system. At last, documents could be produced for universal, machine independent consumption.

There is a new language being developed that will do for applications what Netscape does for documents -- Java. *Java allows transmission of executable programs to any computer connected to the network to be interpreted and played safely and securely in real-time.*<sup>7</sup> Not only will Java threaten Microsoft's empire, it will significantly affect the power needed in a home terminal. Through Java, software programs become network-based and terminal independent. No longer will you be limited by either the application programs or files resident on your home computer -- with it you could leverage the power and information available from the millions of computers attached to the network. As the power of the network goes up as the square of the number of computers connected to it, such an arrangement should have enough power, well, to balance my cheque book, anyway.

This leads to the concept of a hollowed-out computer -- terminals with only enough smarts to efficiently interact with the network. Various referred to as network appliances, network computers, and surfing machines, commercial units will be available by Christmas time.<sup>8</sup> Sounds portable -- sounds ideally suited for wireless networks. Sounds like a security nightmare.

Though the migration of intelligence back into the network may sound like a return to the mainframe era, it will not be. Unlike in the days of the mainframe, surfing machines will not be 'dumb' and network intelligence will not be centralized. Surfing machines will likely be high-end processors optimized for executing code obtained from the network and network intelligence will distributed over potentially thousands of machines. Should these developments pan out, they will lead to computer/telecommunication integration on a scale not even imagined a few years ago. Given that viruses represent executable code often obtained from the network, the proliferation of devices specifically designed to both obtain

and execute code delivered from the network should give the virus business some significant growth opportunities.

As mainstream software development is still largely an art, it is not yet possible to properly verify software. Until it is (if ever), exploitable bugs will remain as will the many unpleasant surprises that often ensue when systems are presented with input conditions that had neither been foreseen nor tested for. As software interactions begin to occur on a global scale and as they are increasingly used to directly control critical systems, the potential for disaster becomes more real. Unlike hardware failures which are to some extent predictable, software failures are not -- either with respect to their size or scope.

## X. THE ELECTRONIC MARKETPLACE

Convergence from a technology perspective does not imply convergence from a service or business perspective. In fact, quite the opposite is happening. While the media independence resulting from digital capability implies that only one multimedia-capable network should be needed, there are far more things one can do with integrated media and networks than with independent ones. When convergence between computers and communication networks is factored in, the number of mainstream and niche opportunities expands significantly.

Considering that in excess of a trillion dollars a day is already transferred electronically and most of this is on private systems at a corporate level, the development of electronic cash and the downmarket migration of these capabilities to the consumer implies that we have not seen anything yet. In the era of electronic cash, there will be numerous currencies available and a resurgence in the art of counterfeiting might be on the way. Overall, the impact of electronic commerce on this scale will be significant -- value chains will be redrawn and businesses will both be created and destroyed in the process.

As more is done on these information superhighways, there will be more to protect -- and more ways to do harm and reek havoc. 'Ladies and Gentlemen, you have your work cut out for you.'

## XI. WE'VE ONLY JUST BEGUN

In a parable concerning the Emperor of China, the Emperor offered the inventor of chess any prize he desired. The inventor opted for a mere grain of rice on the first square of the chessboard, two grains on the second, four on the third, and so on. While the Emperor thought the request modest, a simple calculation shows that at square sixty-four, this would amount to enough rice to bury the entire planet. George Gilder points out that in terms of progress in electronics, we are on square thirty-two.<sup>9</sup> Of course, while the Emperor could solve his problem by beheading the inventor, the solutions to our information age problems are not so straightforward.

Actually, we have no way of even knowing how big the 'chess' board is, if in fact, it is limited at all.

In terms of interesting times, we have only just begun.

#### REFERENCES

- [1] R.P. Weber & R.F. Mortenson, "AT&T's Plan 2000: A Future View of Business Communications", NEC's *InfoVision: Visions of the Information Age*, Aug. 1993.
- [2] V. Schnee & A. Tumolillo; **The End of the RBOCs**, Probe Research Inc., 1994.
- [3] C. Anderson, *The accidental superhighway*, The Economist, July 1, 1995.
- [4] C. Dobson, *Future telecommunication directions and implications*, Wescanex '95, IEEE, May 1995.
- [5] P. White, *The changing role of switching systems in the telecommunication's network*, Communications Magazine, IEEE, January , 1993.
- [6] C. Dobson, *Technology, education, and the information highway*, Learning in the Fast Lane, Edmonton Public School Board, November, 1995.
- [7] G. Gilder, *The coming software shift*, Forbes ASAP, August, 1995.
- [8] N. Gross, et al, *Internet Lite: who needs a PC?*, Business Week, November 13, 1995.
- [9] G. Gilder, *Keynote address*, Telecon '95, CBTA, September, 1995.

**Craig Dobson** is a Senior Consultant in AGT's Technology and New Market Development area. He is chair of the Natural Sciences and Engineering Research Council of Canada's Grant Selection Committee on Communications, Computers, and Components Engineering. Craig is a Professional Engineer and holds an M.Sc.(Eng.) in Electrical Engineering from Queen's University in Kingston.