Electronic Commerce on the Internet

TABLE OF CONTENTS

- Return to JLI's published papers page
- Return to JLI's home page

Electronic Commerce On The Internet - section 1

Andy Johnson-Laird & William R. Trost++

INTRODUCTION

What Is Cyberspace?

Throughout this paper there are references, directly or indirectly, to something called "cyberspace". Cyberspace is one of those widely used terms that has no accurate definition - it is that place not in a computer, or in a computer network, but is that twilight zone in-between computers. The definition is more easily inferred than implied: Where is the money in a person's checking account? Is it in a cash drawer or a bank bag somewhere? Does it really exist? Is it just a data record in a computer somewhere? Does it exist as a data record on more than one computer?

The money exists in cyberspace. It can be removed from cyberspace and brought into the real world by a brief interaction with an automated teller machine.

Return to Table of Contents

WHAT IS THE INTERNET?

The Internet is an alien world and it is difficult to describe a whole world in just a few sentences 1. Perhaps the easiest place to start is with a gentle introduction to the underlying technology that makes the Internet the unique world that it is.

Take Two Computers

The Internet started with humble beginnings. Imagine two computers, A and B, connected together with a data cable. Over this data cable, the two computers can, with appropriate software, transmit the contents of data files either from the disk in computer A to the disk in computer B, or from B's disk to A's. These data files might contain electronic messages, word processing documents, diagrams, sounds, still video images, or full motion video - albeit the sounds and images are in a digital form that can be stored on computers.

Although this scenario suggests that there are human computer operators involved in arranging for such a data file transfer, it can easily be done

without human intervention. This is especially true for more advanced workstations and mini- and main-frame computers. In fact, it was two main-frame computers that were used in 1969 for the first experiments that led to the Internet.

The room-sized main-frame computers of yesteryear, apart from their size, have several characteristics that set them apart from personal computers. First, they are left running all the time; second, they are sufficiently powerful that they permit several users, each on his or her own terminal, to use the computer simultaneously; and third, because they support multiple users, it is quite common for them to provide electronic messaging between those users who might be scattered throughout one or more buildings.

Take More Computers

Most people, on seeing a wire running between two computers, will use a mental model of the two computers "talking" to each other - rather like two people talking on the telephone - first one person says something, then the other responds, and so on.

In fact, the analogy is more like that of a submarine telephone cable or satellite communication link supporting hundreds, if not thousands, of simultaneous conversations. Two computers connected via a wire can, by skillfully designed software, create the illusion of multiple simultaneous "conversations" by slicing each large chunk of information into smaller "packets", ensuring (like a real-world packet in the postal system) that each packet has the recipient's and sender's addresses, and then sending it on its way. By splitting large chunks of information into smaller packets as they are being sent, intermingling these packets with those from other chunks of information while in transit, and reassembling the packets in their proper order on receipt, more than one "conversation" can appear to travel down the connecting wire.

Imagine standing on a bridge over a two-lane highway; looking down you would see many different cars and trucks passing by underneath. Imagine also, a large group of people traveling together, but each in his or her own car. Would they all need to be in adjacent "line astern" formation in the same lane? Given that all of the drivers knew their common destination point, it

would not matter at all if other vehicles interposed themselves - each driver would make the necessary decisions to arrive at the destination, albeit with some variation in arrival time. The drivers would take the appropriate detours should road work or bad weather make some roads impassable. In fact, extending this mental model slightly, the drivers could even travel by different routes, one preferring a scenic route, another preferring to travel on the freeway, and yet all could arrive at the appointed destination.

If all the computers in the network were provided with information about all the other computers in the network, then the actual routing of data packets from Portland, Oregon to, say, Princeton, New Jersey, could be chosen to avoid "traffic jams" along the way. Geographical distance becomes irrelevant; routes can be chosen that are faster, not necessarily shorter. A classic example of this is the case of one of the authors in Portland, wanting to send a message to the other author in Beaverton, eleven miles to the west of Portland - the data packets in question travel to Stockton, California, then to Washington, D.C.; on to New York; then to Hartford, Connecticut. From there, they go to Cleveland; then Chicago; back to San Francisco; up to Seattle; over to Spokane, and finally, over to Beaverton, Oregon.

While this 6,000 mile cross-continental journey may seem like crazy meandering to deliver a message eleven miles, 2 one must bear in mind that the "super-highway" simply does not go from Portland to Beaverton, and therefore, the quickest electronic way might not be the shortest. Furthermore, Sprint handles Portland traffic, whereas GTE handles Beaverton. The miracle is that it happens at all. The routing is merely a technical convenience that serves to illustrate the prophecy in Gertrude Stein's comment (albeit of Oakland, California) that there is no "there" there on the Internet - everything is "here" 3.

A side-benefit of this massive communications web is that there is massive redundancy in the system. Should one part of the network fail, packets can be re-routed out of harm's way. In fact, the original purpose of the so-called packet switched network was to provide redundancy in case a Russian missile took out, say, Washington, D.C. 4 For reasons that history has hopefully rendered obsolete, the world, especially the U.S., has a global communications web with massive built-in redundancy. But what a web!

Take Three Or Four Million Computers

By current estimates there are more than 3.2 million "host" computers connected to the Internet, and they in turn provide access to dozens of other computers in their immediate vicinity. These computers vary in the immediacy of their connection to the Internet in the same way that some people live closer to a freeway on-ramp than others.

At the bottom of the connectivity ranking are personal computers that only dial into, say, CompuServe, whenever their owners choose, be that once a day or once every other week. Other computers, typically desktop workstations running operating systems such as Microsoft Windows NT or UNIX (which is really a minicomputer operating system scaled down to the desktop), will dial into computers that are more closely connected to the Internet. In this case, their "contact points" with the Internet are on an automated schedule and occur perhaps two or three times every hour. Next in the connectivity chain come computers that have a continuous connection with the Internet, albeit via a normal phone line that can transmit data at modest rates of 1,000 characters per second. At the top of the connectivity scale are computers connected to the "backbone" of the Internet via high speed telephone lines capable of transmitting and receiving huge amounts of information hundreds of times faster than an ordinary domestic voice line.

We no longer marvel at the global telephone network - but we should. The ability to pick up a telephone in almost any country in the world, dial a phone number in any other country, and be connected within 20 seconds is nothing short of miraculous.

The one technical leap that makes the Internet perhaps the largest construction ever created by our species is the ability for millions of computers to "pick up the phone" and talk to each other almost simultaneously - or at least in such rapid succession that we plodding humans perceive it as near simultaneous. By providing that "near simultaneous" connectivity we have perhaps made that same kind of leap that neurons in a primitive brain made when homo sapiens first became conscious of its own existence. While this author certainly does not profess to understand the dividing line between the self-conscious brain and the non-self-conscious brain, it appears clear that some "magic" happens merely

because of the number of neurons in the brain and the massive interconnectivity of those neurons. That same magic appears imminent in the Internet, as can be seen by the effects of scale applied to the various capabilities of the Internet. Having a massive network of computers, each capable of communicating with each other or, on a whimsy, communicating with groups of other computers with no regard to geographical distance, has created a shimmering web of global intercommunication that simply has not happened before in humankind's existence.

Return to Table of Contents

Geography No Longer Exists

Almost as a byproduct, the Internet has achieved something that hitherto only science fiction writers could contemplate - it has removed geography. Distance has no significance. City, county, state and country borders are meaningless, for only "atoms stop at borders, electrons do not"[5].

This electronic abolition of geography will cause the largest upheaval in the domain of the law since its inception - the law depends upon latitude and longitude to determine which law applies, and this underpinning has been removed, although most legislators do not yet understand this fact.

Return to Table of Contents

The Changing Face of Commerce

The Internet is also about to change global commerce in a manner that has never before been seen. Individuals in the privacy of their upstairs bedroom can, in a matter of a few seconds, be global purchasers and exporters of digital information (be it text, images, video or audio), purchasing directly from vendors and selling directly to purchasers, without the need for wholesalers, retailers, customs brokers, shippers, and all those who make their living moving "atoms" around the world.

Whilst this might seem futuristic, it is, in fact, the present. All of the technology is in place right now; people are conducting commerce on the Internet right now.

But commerce on the Internet is very much like a teenager's view of sex: One thinks that everyone else is doing it, but there are not that many who really are doing it, and for those who are, it is a clumsy experience. However, with people and businesses arriving on the net in great numbers, one can see with reasonable certainty that, now that the techno-dweebs have built it, the rest of the world will surely come.

Return to Table of Contents

Am I Me? Are You You?

Business on the net faces several fundamental problems that have been solved in the real world: (a) am I really me? (b) are you really you? (c) is the merchandise that you say you have for sale, all that you say it is? (d) did I really order it? (e) can I really pay for it? and (f) whose law (if any) applies to the sale?

Of these questions, only (a), (b) and (d) are technological problems concerning identity and authentication - the others are questions that must be answered in isolation of the computer technology (but just get harder to answer with computer- mediated commerce).

Identity On The Net

The question of identity in the real world has been solved (at least to some practical level) by the creation of a web of trust. You can trust that I am me because I carry pieces of paper and plastic that were given to me by agencies that you can trust would not have done so if they thought I was not who I say I am. This web of trust is not fool-proof (as many a fool has proved) - it is all too easy to purchase social security cards, credit cards, green cards, and bogus passports, but, by and large, it works, and most people are who they say they are, most of the time.

In more arduous circumstances, you know that I am me because I can, without too much hesitation and on demand, make unique squiggles on paper. Someone else (otherwise known as a notary public) is prepared to say that he or she has checked the aforementioned pieces of paper and plastic,

and verified that my most recent squiggle still looks like my earlier squiggles, and the ugliest photographs taken of me since childhoold are still somewhat similar to how I appear in the flesh.

In the world of electronic mail and typewritten messages this web of trust does not yet exist to any practical level. An electronic mail message that says:

- **This is President William Jefferson Clinton. I would like to buy a copy of your software. Please send it to me at 850 NW Summit Avenue, Portland, OR 97210.**

will doubtless be greeted with a great deal of incredulity because it fails obvious "sanity" checks. But what if the name were changed to my own? How could you as a vendor, receiving the message from me, have any certitude that it was really me who sent it and not someone merely posing as me?

Worse yet, what happens if, having received the merchandise, I were to deny that I ever ordered it, and refuse to pay for it? How could you prove that it really was me that ordered it?

The Electronic Web Of Trust

The only viable basis for establishing identity is by the use of a certification agency - that is, a neutral party who can be trusted to vouch for individuals, and who, as in the real world, can issue unique tokens to certify that individual Internet users are who they say they are.

As in the real world, this web of trust woven by certification agencies will be built upon the foundational concept that certain physical tokens (driving licences, passports, etc.) are sufficiently hard to obtain that most of them are issued to the right individuals. As mentioned above, this web will not be unbreakable.

As society comes to realize the importance of establishing an individual's identity for electronic commerce, doubtless more emphasis will be placed on ascribing unique identification to the newborn around the world. The price

of certain identity is paid for in the currency of individual freedom, and doubtless, overzealous governments will attempt to solve the problem with draconian means ranging from UPC bar codes tattooed on body parts to biometric identification of unknown validity 6.

Real Identity Versus Reputation?

For commerce on the Internet it could be argued that one is less interested in identity than reputation. Other than for market research reasons, a merchant need care little more about a purchaser than the level of certainty with which the purchaser will pay their bill. The purchaser, on the other hand, will likely care little more than the merchant's ability to deliver the merchandise in working order, at the agreed price, and in a timely way.

This is essentially the situation that exists in today's mail order business. The purchaser quotes a shipping name and address, and cites a credit card number and expiry date, and armed only with this and knowledge of the purchaser's recent reputation from the credit card issuer, the merchant ships the merchandise.

The sad truth is that plastic credit cards afford plastic identity, as any merchant and many purchasers can confirm. It is all too easy for malfeasants to misappropriate credit card numbers and, for a short time at least, pass themselves off as the legitimate purchaser and run up enormous bills.

Public Key Encryption Is An Answer

The fundamental underpinning for electronic commerce is for you to know with some level of certainty that I am who I say I am, and that computer-based documents purporting to emanate from me really do. Once that is established, my reputation can be determined by inference from other existing means.

To this end, public key encryption promises to be a useful tool. With an element of historic irony, public key encryption works like one of the oldest legal documents, the indenture.

An identure consists of a legal agreement that was handwritten in side-by-

side duplicate on a single piece of paper. The two copies of the agreement were then separated by cutting the single piece of paper into two along an arbitrarily wavy line. In the future, if there was any claim that one document was bogus, it could be authenticated by placing one copy of the agreement alongside the other and demonstrating that the "indentured" randomly cut line matched the other part exactly.

A public key works in a similar fashion as can be seen from a simple example. Imagine that you wish to use public key encryption - the software you need is freely available at no charge throughout the world via the Internet 7. Known as PGP (Pretty Good Privacy), this software will run on almost all personal computers in the world today.

Using PGP, the first step is to generate your own private and public keys. These two keys, like the indenture above, are complimentary but, unlike the documents written on the indenture, are different. The diagram below shows the generation of your private and public keys - a process that involves the generation of some random numbers taken by timing random keystrokes on the keyboard and the selection of a (hopefully) unguessable pass phrase. The pass phrase is like a password but can be more than a single word, and should be selected by the user based on some obscure fact or phrase (preferrably a meaningless but memorable phrase) unlikely to be known by others (for example, my pass phrase is a phrase I saw on the notice board in my brother's kitchen some 30 years ago - not even he remembered it). The key generation phase is shown in the following diagram:

Generation of private and public keys

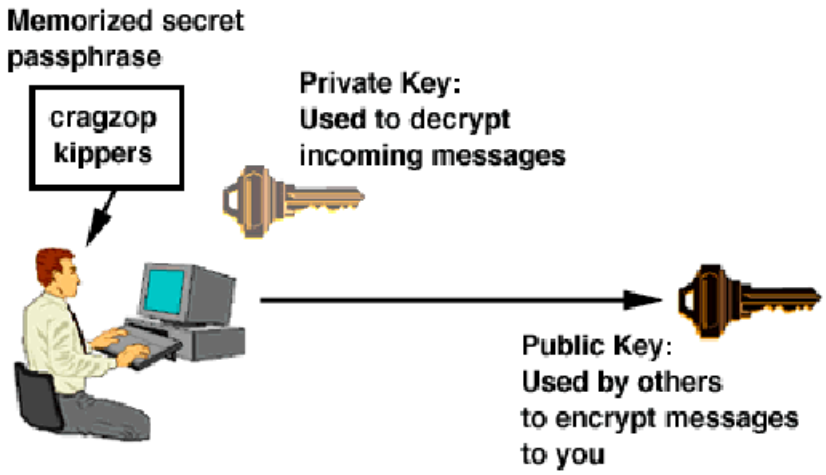

Memorized secret passphrase (in this case: cragzop kippers) is used to decrypt incoming messages in conjunction with the Private Key.

The two keys live up to their names: to keep messages to you confidential, the private key must be kept under your control, and the public key is made public by sending it to anyone who might wish to transmit confidential email messages to you. For example, I transmit my public key as part of the "signature" lines I append to electronic mail messages:

```
+=================================================================+
| Andy Johnson-Laird | Forensic Software/Multimedia/Image Analysis |
| Johnson-Laird Inc. | Software Engineering/Reverse Engineering    |
| 850 NW Summit Ave. | Software Plagiarism Assessment              |
| Portland OR 97210  | Software Project Failure Analysis           |
+-----------------------------------------------------------------+
Tel: (503) 274-0784  | FAX: (503) 274-0512 | andy@jli.com

For encrypted messages use the following public key:

-----BEGIN PGP PUBLIC KEY BLOCK-----

Version: 2.6
mQBtAy8z3f4AAAEDAMScuLJ1Du/Vuhb4kqHOX4TKqODOvUwoqQAm3bufdVSlpwo4
VH6uTpmIbdQBHFQ+Gvj6NRVF0KmACI0I2pA2Pwfhgrx7b5L4XOTyjZiM92QfmpyN
IcuH65DUiYqbhU8x3QAFEbQsQW5keSBKb2huc29uLUxhaXJkIDxhbmR5QGpsaS5w
b3J0bGFuZC5vci51cz6JAJUDBRAvNc0ieKLFbnNuvD0BAdCvA/4w8896w7QodgKL
KKsKb+9xRwShy1/ygkbmepungnky4VhM6rJX+JdPgSh/uSAaQk4Wt6UnUuVnMy1P
HSVNStOVLFzINg2WMB8OdCi+OGjQ8r03AtsHqXsnWMvF3t6kpOJ7KaZuCg8LvLUY
nk7u+YPJk7yV/MOJbl6NS7nSYYa/+IkAdQMFEC8z3xDUiYqbhU8x3QEBOroDAKiN
EC18JyEoubuqKvfIdehNSJZ4cJCC+AV695o1lbq1ww3PUW5R2jWuyd7Nx5RZOu6n
/V3dkApvSxBG/rLPiBrsE36cz6CWTeWhIQdOhendqSREELGtTMKkvqpbDKWIJw==
=Yku8

-----END PGP PUBLIC KEY BLOCK-----
```

The PGP public key is the mass of apparently gibberish characters at the end of the signature lines. When these textual characters are viewed in the binary zeroes and ones that computers use, the gibberish is a long binary number which is the actual digital key.

The application of this technology to business on the Internet, is twofold:

- 1. It can be used to encrypt credit card information for transmission to merchants (or any other confidential aspects of a transaction).

  2. It can be used to authenticate merchandise orders by generating a digital signature.

Anyone who receives this public key can use it with the PGP software to take a data file (which could be plain text, a, WordPerfect document, or WordPerfect itself, or anything else) and can encrypt that file and transmit it electronically to me. Once I receive it, using my private key (which I need to unlock for use with my pass phrase), I can then decrypt this incoming file. This process is shown in the following diagram:

Receiving an encrypted message

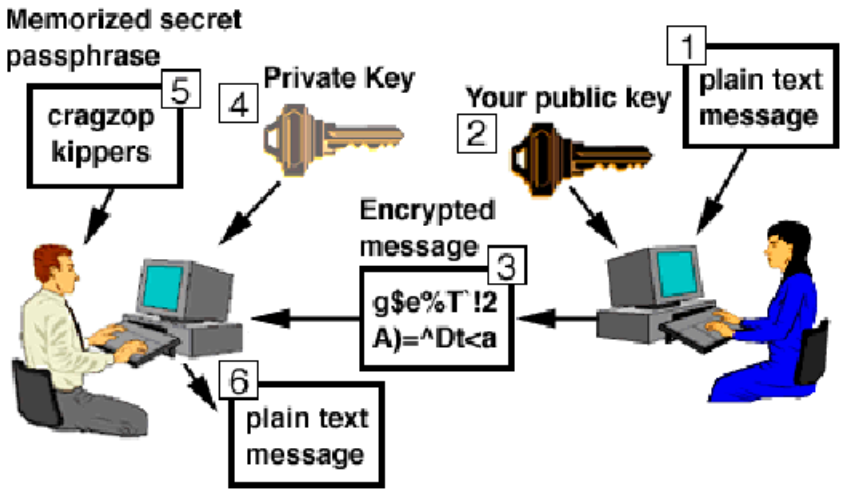

Third party wants to send you a confidential ecrypted message.

The security of the encryption depends upon the number of binary digits that are used in the encryption key. Using 1,024 binary digits results in "military grade" encryption.

Digitally signing a document and verifying its signature is a similar process, but the order in which the public and private keys is reversed, as shown below:

Sending an encrypted message

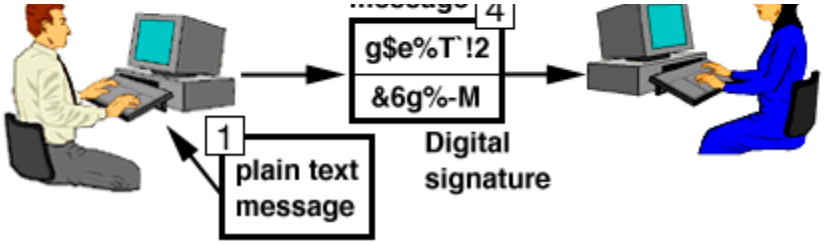

Third party authenticates contents of document and that it was you who sent it.

A digital signature attached to a document permits the recipient to do two things:

- 1. To verify that the message received is undamaged - the contents of the message is interwoven with the digital signature and there will be a mismatch of signature to contents if the message has been tampered with.

  2. To authenticate that the owner of a particular public key really was the sender of the message.

It is these two capabilities that will make PGP and other public key systems particularly useful for business on the Internet, although, to be fair, it will also shift a thief's attention to stealing private keys and pass-phrases. The heart of PGP technology is to be able to trust the public keys which one receives from others, and this brings us squarely back to contemplation of certification agencies - organizations that the general populace can trust as having made a reasonable effort to correlate a given public key with its rightful owner.

PGP already has the seeds for a web of trust insofar as it permits public keys to be "signed" digitally by third parties who have authenticated (or have every reason to trust) that the owner of a public key is the person they say they are. The problem, of course, is that you are unlikely to be filled with quiet confidence if you see that my key has been signed by 38 people that you do not know. The situation is therefore ripe for a government agency, perhaps the U.S. Postal Service, to act as a public key repository.

Return to Table of Contents

Digital Cash 8

The major problem with the use of credit cards on the Internet, apart from the exposure to public gaze if one transmits the numbers in clear text, is that the merchant and the bank (and anyone else with access to the bank's records) knows who you are, what you bought, when you bought it, from whom you bought it, and for how much. To those who have grown up in the age of the credit card and the resulting swath of credit history we each leave behind us, this may not seem particularly distasteful, but modern society places great stock in being able to preserve some vestiges of personal privacy from the gaze of Big Brother and all the Little Brothers with their camcorders and computers.

Digital cash, though still in its infancy, offers the prospect of being able to conduct business on the Internet and yet retain personal privacy to a degree determined by the individual, not by the merchants or the banks. It also represents a first-time opportunity for individuals to conduct global electronic commerce in electrons while conveniently ignoring sales tax, value added tax, use tax, and any other tax which governments can invent.

Digicash, a company in the Netherlands, has been running experimental systems over the past few years, using a cleverly constructed system for purchases and payments 9.

Digital cash is already in use on a limited scale, although it is still sufficiently young that, as far as this author is aware, there are no recorded cases of theft of digital cash. The speed of its adoption remains limited by the willingness of society to move from folding and plastic money to truly digital money.

Return to Table of Contents

The Internet Is Global

An individual with a personal computer in his or her back bedroom can reach out over the Internet to buy or sell around the world, dealing directly with vendors or purchasers of digital information without the need for

middle-men - no advertising agencies, buyers, wholesalers, retailers, shipping clerks, customs brokers, customs officers, and UPS/Fedex deliverymen are required. If one has a fixed or low price access to the Internet, the distribution costs of digital information becomes vanishingly small - the Internet, in effect, is acting as a frictionless, free, global distribution pipeline. The ease with which this business can be transacted is of clear benefit to the individual; it also permits them to bypass, deliberately or accidentally, import/export regulations and customs tariffs.

Furthermore, it remains to be seen which bodies of law (if any) will apply to the individual in Hawaii, who orders digital information from a vendor in Taipei, makes digital payment to a digital bank in Switzerland, and receives the information from a computer in the Grand Caymans 10.

Clearly, there will need to be global laws to deal with this new form of global commerce, protecting the merchants and purchasers alike. One can only imagine how the governments and the middle-men of the world will react when they realize they have already been squeezed out of the deals. Choice of venue for any litigation arising out of such deals will also, as the Chinese curse 11 describes, be "interesting."

Return to Table of Contents

About the Authors

Andrew Johnson-Laird is a forensic software analyst and president of Johnson-Laird Inc., a consulting company specializing in Techno-Archeology (the study of failed software projects), forensic software analysis, software plagiarism assessment, Internet policing for stolen software, and "clean room" software reverse engineering. He has thirty years' experience of programming, managing programmers, and developing interoperable and competitive computer software.

William R. Trost is a consultant to Johnson-Laird Inc. and specializes in computer system software and Internet consulting, both in the forensic and general commercial contexts. He acts as system administrator for Johnson-Laird Inc.'s workstations and beats the "user-vicious" Unix operating system

into submission on a regular basis.

Return to Table of Contents

- Return to JLI's published papers page
- Return to JLI's home page