Internet Commerce: Evolution, Examples and Evaluation

L. Jean Camp
Department of Engineering & Public Policy
Carnegie Mellon University
Pittsburgh, PA 15213
camp+@cmu.edu

Prepared for Banking on the New Paradigm
Tokyo
December 1995

1. Introduction

Electronic commerce has arrived. Consumers today use multiple electronic commerce protocols: phone banking, ATM transactions, debit and credit card purchases for both remote and point of sale purchases. Electronic billing, automatic deposit and withdrawal are ubiquitous.

New forms of electronic commerce include commerce with trusted hardware and Internet commerce. In this work I focus on Internet commerce. With Internet commerce goods are advertised, purchased and sometimes delivered on-line.

In order to explore Internet commerce I begin by attempting to define the Internet. As there are nearly as many definitions as there are users, I approach this problem by offering high level definitions of the underlying protocols, and a brief history.

Having defined the environment, I consider specific Internet commerce protocols: Digicash (Chaum, 1985; Chaum 1988; Chaum, 1992), NetBill (Cox, 1995; Sirbu, 1995), First Virtual (First Virtual, 1995) and the Secure Electronic Payment Protocol (Mastercard, 1995). For each protocol I discuss a specific transaction, including the information fields and encryption operations. I emphasize the implications of the technical aspects. Finally, I compare the selected systems and make some general observations about Internet commerce.

2. The Internet

2.1. Underlying Technology

The fundamental technology of the Internet is the IP protocol. IP provides the delivery of data.

IP is a *connectionless* protocol. This means that each data packet is delivered independently, so that differing networks can communicate. In contrast, telephone networks have traditionally been connection oriented. Connection oriented protocols establish a point to point connection, from one phone to another, when communication is requested. This is simple in a homogenous environment, but difficult when the connection must pass through heterogeneous networks. The Internet protocol provides the ability to transport data packets through a network of networks. This enables, for example, DEC stations using Ethernet to communicate with DOS machines on LANs.

Much information cannot be transmitted in a single packet. The *transmission control protocol,* TCP, accepts discrete packets and orders them to provide continuous data flow. TCP provides flow control, sequencing and error detection. TCP provides the orderly and reliable delivery of data.

By 1991 the TCP/IP protocol suite consisted of about one hundred protocols. By 1991 there were more than 700,000 machines using TCP/IP to connect 4,000,000 users. (Cerf, 1993)

Upper level protocols, such as telnet and the *file transfer protocol* (FTP), use TCP/IP. FTP was an early protocol that allowed users to 'publish' documents to the Internet community by making them available for retrieval using a simple command line interface.

The *hypertext transfer protocol* (HTTP) is an application that provides seamless delivery of different types of data with a user-friendly graphical interface. HTTP is the protocol of the World Wide Web. HTTP allows users to easily publish, locate, and obtain information on the Internet. It provides a simple user interface which highlights other files using color or graphics. HTTP catalogs locally available applications for file display, and automatically provides the selected text, sound or graphic using these local applications.

Internet commerce has increasingly become possible with the advent of the World Wide Web. The Web is growing at many times the rate of overall Internet host growth. The Web allows the consumer to locate information of interest on the Internet without requiring any technical expertise.

All Internet commerce proposals can be used with the Web. In addition, some commerce protocols (Mastercard, 1995; VISA, 1995) are comprehensive and include the ability to transfer funds using only email. (Email requires only an application to compose and read email and TCP/IP).

For a detailed discussion of network protocols see (Schwartz, 1987) and (National Center for Supercomputing Applications, 1995).

2.2. Interoperability and Nonrepudiation

Although Internet communications protocols are interoperable, Internet commerce systems are not. A consumer cannot cash a NetCheque and obtain Digicash. Internet commerce systems which use open standards may become interoperable. Proprietary standards create barriers to interoperability. Proprietary cryptographic standards are particularly problematic, since interoperability with proprietary standards requires alternative Internet commerce providers to place blind trust in competitors' decisions.

Interoperability requires some common standards and, especially in the case of commerce, shared trust. In order for there to be trust, electronic signatures and dates must be verifiable. The Financial Service Technology Consortium (FSTC) is working toward standards for interoperability electronic commerce (FSTC, 1995).

2.2.1. Digital Signatures

Commerce systems require that there be proof of payment or promise to pay. This means that reliable commerce system require *nonrepudiation.* Digital signatures create nonrepudiation, meaning that holder of a document with a digital signature can prove that it was constructed by the signer.

Recall from (Camp, 1995b) that encryption with a symmetric key assures only that one of the parties to a communication created a document, and cannot specify which party did so. Thus only encryption with the private key of a set of public keys provides nonrepudiation. (Also recall that encryption with the published key of a set of private key assures that only the intended recipient can read a message.) An extensive description of cryptographic tools and systems can be found in (Schneier, 1995).

Digital signatures depend on two factors: the security of the secret key and the strength of the binding of that key to a physical entity. If the security of the secret key is lost the creation of forgeries is trivial. Legal issues of responsibility for a lost key and liability for resulting losses have not been determined.

A physical entity, such as a merchant, can be bound to a public key in three ways. First, a trusted entity, such as a bank, may verify the key upon request. This requires that the central trusted authority be highly available.

Second, the merchant can have a document signed by a trusted authority that contains both identity information and the public key. Such a signed document is called a *certificate*, and there are open standards for certificate formats.

A third technique, which is used in the Pretty Good Privacy system, requires having a certificate that is signed by many partially trusted authorities. The binding of a key to an identity is strengthened with an increase in the number of partially trusted signatures. The use of multiple signatures means that the binding could remain valid if the security of the key of one signer is compromised. The disadvantage is that verification of a certificate is more expensive, since multiple digital signatures must be checked.

Issues in certification of public keys include operational questions, such as the *lifetime* of the certificate, and the avoidance of bottlenecks without too much distribution of authority. The period a certificate is valid is the lifetime of that certificate. A longer certificate lifetime results in a greater possibility of fraud. A shorter certificate lifetime results in a greater load on the trusted authority.

Verisign is attempting to serve the market's need for a trusted authority. Verisign provides authentication of identity for Netscape, Open Market, IBM, Internet Factory, the Internet Office Web Server, the WebSite Professional server, and StarNine's SSL Security Tool Kit (Verisign, 1995). Verisign provides varying levels of authentication. An individual can obtain the lowest level by simply claiming a public key, in which case Verisign only verifies that this person claims to hold that key. The highest level of certification requires some physical proof of identity, such as a passport.

While Verisign is focusing on merchants, the United States Postal Service is attempting to serve the market for consumer public key verification. Under the United States Postal Service plan a consumer would bring some proof of identify, such as a passport, and a disk containing a public key to the Post Office. Then the key holder would swear under oath to a Postal Employee to be the individual claimed. The employee would then provide the key holder with a signed certificate on the key holder's disk. Since the individual swears to a Federal employee, misrepresentation in this case is a federal crime.

2.2.2. Digital Time Stamps

In addition to being able to verify the originator of a document, it is often necessary to verify the date the document was signed. *Time stamps* prove that a document was signed at a particular time. In some Internet protocols, time stamps are provided by including an agreed-upon time in the document to be digitally signed by all concerned. Surety Technologies is providing reliable time stamps for general use. (More information is available at http://www.surety.com/about-surety.html.)

Surety uses a patented chaining technique (Haber, 1991). To verify a document, a hash value is sent to Surety. The hash of this document is combined with other documents to be dated in a binary tree, producing a final value that is widely published. Currently the final, or root, values are printed in the New York Times classified advertisements. Surety returns the hash value of the original document and the other hash values necessary to construct the final value. The series of hashes can show that document B was signed before document C and after document A. Using this technique, Surety can identify a document as being signed in a particular hour on a given day. In order to forge a date, it would be necessary to change every document signed after the falsified document on a given day, or every copy of the New York Times.

2.3. Internet History

The Internet began as the ARPANET, a United States government project for connecting scientific research sites. The tools for internetworking computers were developed by scientists and researchers for use in their own nonhierarchical heterogeneous computing environments. The techniques developed were designed for distributed support, with comments from the entire community sought and considered. Although the ARPANET consisted of only a couple of hundred of computers at that time, it created the core of compatible inter-networked computers that became the Internet.

By 1983, all the networks connected to the ARPANET used TCP/IP for communication. After the release of Berkeley Unix 4.2, TCP/IP was included in every Unix workstation. The Unix standard created a commercial opportunity for network products (Cerf, 1993). Although the vast majority of these machines were not initially connected to what we now know as the Internet, the ability to network networks became a standard feature for high-end operating systems.

In 1986 ARPANET became NSFNET, and its mission expanded to include students and libraries as well as researchers. In 1990 the first commercial email provider, MCI Mail, was connected to NSFNET. Along with commercial email providers, commercial information providers came onto the Internet. Early adopters of Internet technology for the sell of information include Dow Jones and Dialog (Cerf, 1993). Thus began Internet commerce.

By 1990 the growth of the Internet was too profitable to be ignored by information providers. However, the market remained primarily technical individuals, with access to information requiring either some understanding of Unix or proprietary software. The growth of the Internet since that time illustrates that the user community has expanded, as shown in Figure 1. (The data used for Figure 1 came from (IDS, 1995a)).

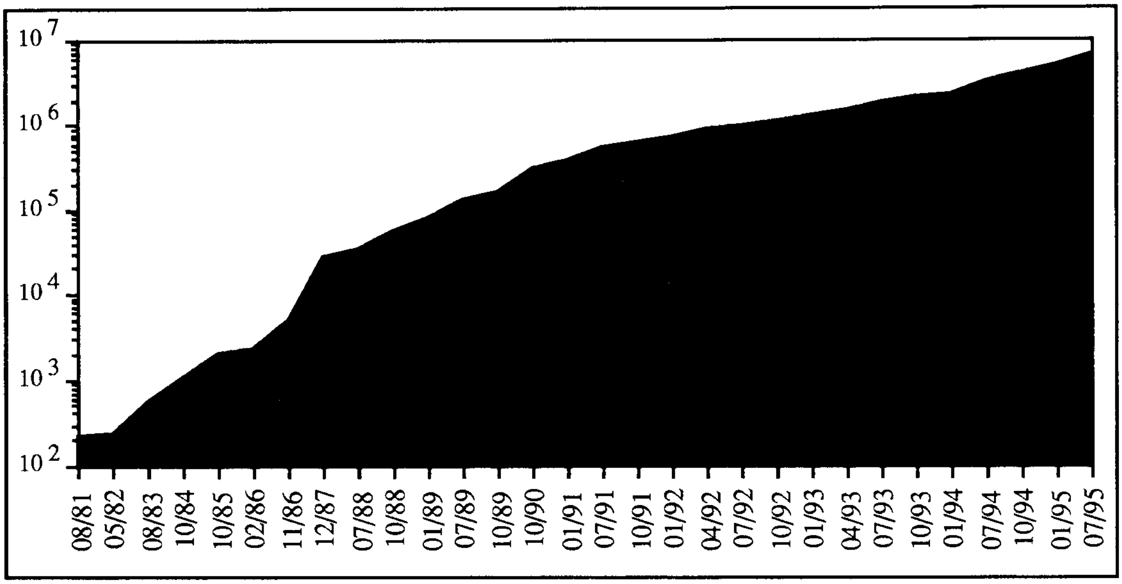

Figure 1: Exponential Growth of the Number of Computers Connected to the Internet

A year before the connection of MCI Mail, a European researcher, Tim Berners-Lee, became concerned with effectively transporting the images, postscript files, ASCII text and data files necessary for collaborative physics throughout Europe. The protocol he developed for collaborative physics is the underlying technology for the World Wide Web. The Web allows consumers to search for a variety of information with a straight forward graphical interface. With the Web, the Internet became fully capable of supporting user-friendly distributed commerce, just as previous protocols had enabled functionality from simple communication to file transmission. Table 1 illustrates how Internet commerce protocols build on previous protocols, which had in turn expanded the pool of possible merchants and consumers.

Protocol	Connects	By Providing
Internet commerce Protocols	Consumer to Merchant	payment, possible delivery verification
Hypertext Transport Protocol	Application to Application	location and presentation
Transmission Control Protocol	Machine to Machine	reliable delivery of multiple packets
Internet Protocol	Network to Network	delivery of packets between networks

Table 1: Hierarchy of Protocols on the Internet

The World Wide Web is a critical element in emerging markets. And, although the Internet began as a specialized US Government project, the Internet is now global. The Internet domain survey has expanded to include ninety countries. The growth of hosts on seven continents from the Internet Domain Survey (IDS, 1995) is shown in Table 2.

The customer base on the Internet grows with number of countries and connections grows.

Region	Hosts in January 94	Hosts in July 94	Hosts in October 94	Hosts in January 95
North America	1,685,715	2,177,396	2,685,929	3,372,551
Europe, West	550,933	730,429	850,993	1,039,192
Europe, East	19,867	27,800	32,951	46,125
Middle East	6,946	8,871	10,383	13,776
Africa	10,951	15,595	21,041	27,130
Asia	81,355	111,278	127,569	151,773
Pacific	113,482	142,353	154,473	192,390

Table 2: Regional Growth on the Internet

2.4. Current Internet Commerce Opportunities

Many successful business ventures are now on the Internet. Table 3 shows examples of businesses on the Internet, and corresponding paper information markets (adapted from Sirbu, 1995).

Market Structure	Electronic Example	Paper Example
Publisher pays	WWW catalogs	Mail order catalogs
Advertiser pays	Lycos, Yahoo	Free weekly papers
Club pays	Claire, Site license software	Corporate library
Customer subscription	Web magazines, dlist	Professional magazines
Customer pay per item	First Virtual	Storefront sales
Customer pay for time	AOL, COMPUSERVE	Rental items
Mixed ads & customer payment	Prodigy, Netscape business sites	Newspaper

Table 3: Structure of Information Markets

The Internet supports a range of business functions, not simply payment. Every transaction has multiple phases: discovery, price negotiation, final selection, payment, delivery, and customer support. The Internet can support all stages of Internet commerce (Sirbu, 1995).

Product discovery is enabled on the Internet through advertising and electronic word of mouth. Products information is disbursed through Web pages, distribution lists and Usenet groups. The Web enables individuals to locate specific information and search by product or company name. Corporate Web sites often exist solely for the purpose of distributing product information with a simple graphical interface. With distribution lists, or dlists, individuals that have a common interest form a closed group and transmit messages of interest to all members of this group. Announcements of new products are made by members of the distribution list. Usually distribution lists are motivated by discussion, with product announcements being a small fraction of the traffic. In Usenet groups new products are announced by subscribers, as is the case with distribution lists. The difference is that Usenet groups are open forums. This implies that not only are product announcements overwhelmed by discussion, but also the information in the groups is notoriously unreliable. Furthermore, direct advertising across Usenet groups is considered offensive by Internet users. Distribution lists, Usenet groups and the Web interact. URL's are sent over distribution list and posted on Usenet, and Web sites connect to archives of Usenet groups and discussion lists. (A URL is a Uniform Resource Locator, i.e. an address for the World Wide Web.)

All of the technologies consumers use to find out about services can also be used to locate suppliers. Web search engines, such as the World Wide Web Worm and Lycos, provide a simple way for consumers with Web browers to locate products.

Price negotiation is supported by email and electronic data interchange. Information goods can be delivered on-line. Customer support can be offered on-line through email and Web pages.

Every phase of a commercial transaction has associated costs. The ability of an Internet commerce protocol to reduce transaction costs depends on its ability to address these costs. For comparison the distribution of costs in a credit card transaction is shown in Figure 2 below (Sirbu, 1995). These cost categories are relevant to the Internet commerce protocols examined in Section 3.

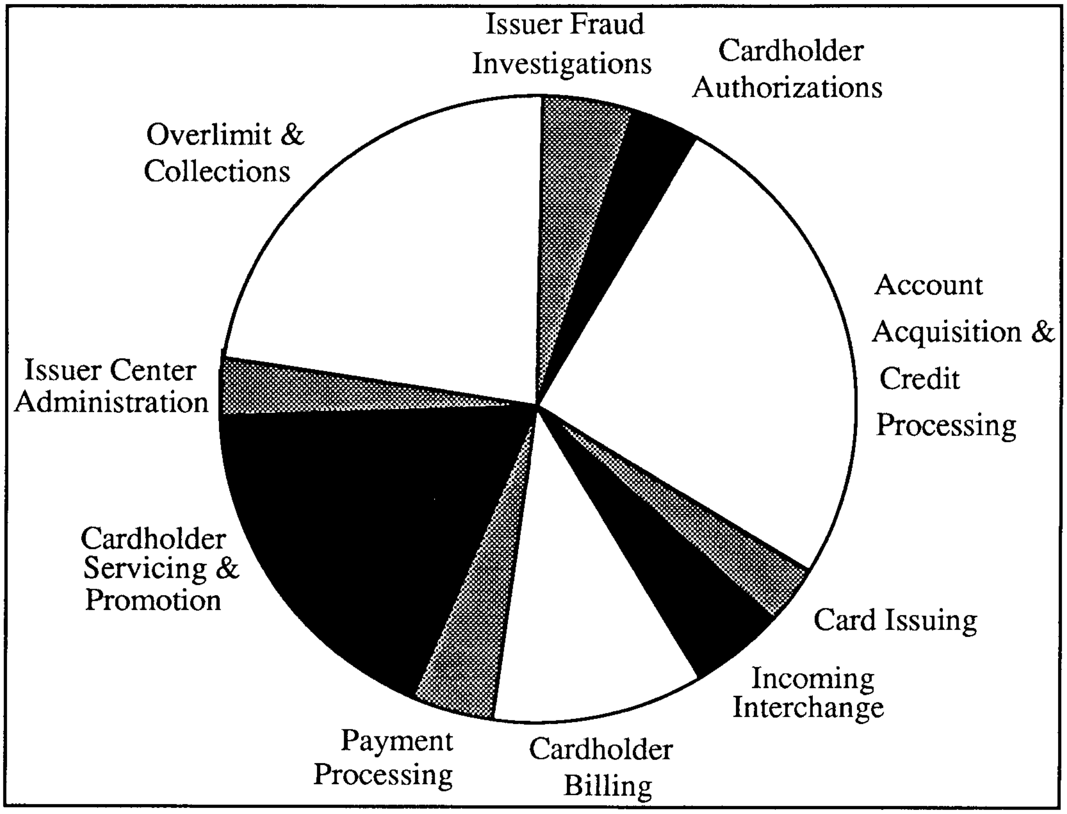

Figure 2: Cost Distribution in a Credit Card Transaction

3. Internet Commerce Protocols

Many of the techniques and assumptions for Internet commerce are not unique to Internet commerce. Internet commerce protocols are multiplying as the Internet expands. Currently one private listing of Internet commerce protocols counts eighteen different Internet commerce protocols and forty places to shop on the Web (Hanushevsky, 1995). This list does not include at least two additional Internet commerce proposals.

In the remainder of this section I will discuss four Internet commerce protocols and illustrate that each has different business assumptions. These protocols are Digicash (Chaum, 1985; Chaum 1988), NetBill (Cox, 1995; Sirbu, 1995), First Virtual (First Virtual, 1995) and the Secure Electronic Payment Protocol (Mastercard, 1995). These protocols provide four different perspectives that cover the range of commerce proposals: token commerce or electronic cash, aggregation of credit or debit purchases with the bank off line, credit or debit purchases without aggregation with the bank off line and credit or debit purchases with the bank on-line.

For each protocol I examine the transaction steps, underlying business model, and security aspects. In the previous paper in this series I offered high level definitions of security, privacy, and basic cryptographic operations which I use here I assume the reader has access to this paper; however, very brief definitions are offered in the appendix. To simplify the transaction descriptions, I use the standard assumptions that the customer is female; the merchant is male; and the bank is a genderless organization.

3.1. Digicash

3.1.1. A Digicash Transaction

Digicash is a token based currency (Camp, 1995a). This means that the string of bits transferred in a Digicash transaction has value. In a notational currency system the information transferred is an instruction to change notations in a ledger, such as a bank's records. In notational currency the value is held in the records, not the instruction. Digicash is intended for both information goods and physical goods.

In the earliest version Digicash offered an anonymous protocol graceful in its simplicity (Chaum, 1985). However, the protocol was not feasible. Tokens could be multiplied to form new valid tokens, that is, consumers and merchants could trivially manufacture cash.

The advantage that the tokens could be verified independently was mitigated by the fact that double spending was perfectly anonymous, and therefore fraudulent parties could not be identified. Finally, there was no verification of payment. Thus, acknowledgment of payment depended entirely upon the goodwill of the merchant (Yee, 1994).

The later version of Digicash (Chaum, 1988) considered here addressed two of these problems. Users can no longer trivially manufacture new tokens. Individuals that double spend can be detected with some probability after the fact. Probability of detection is a function of the size of the fraudulent purchase. Consumers can verify that they have paid merchants; however, this verification requires the loss of anonymity. Merchant fraud remains a potential problem.

Digicash transactions require three parties: a bank, a merchant and a customer. The steps in a Digicash transaction are shown below in Figure 3.

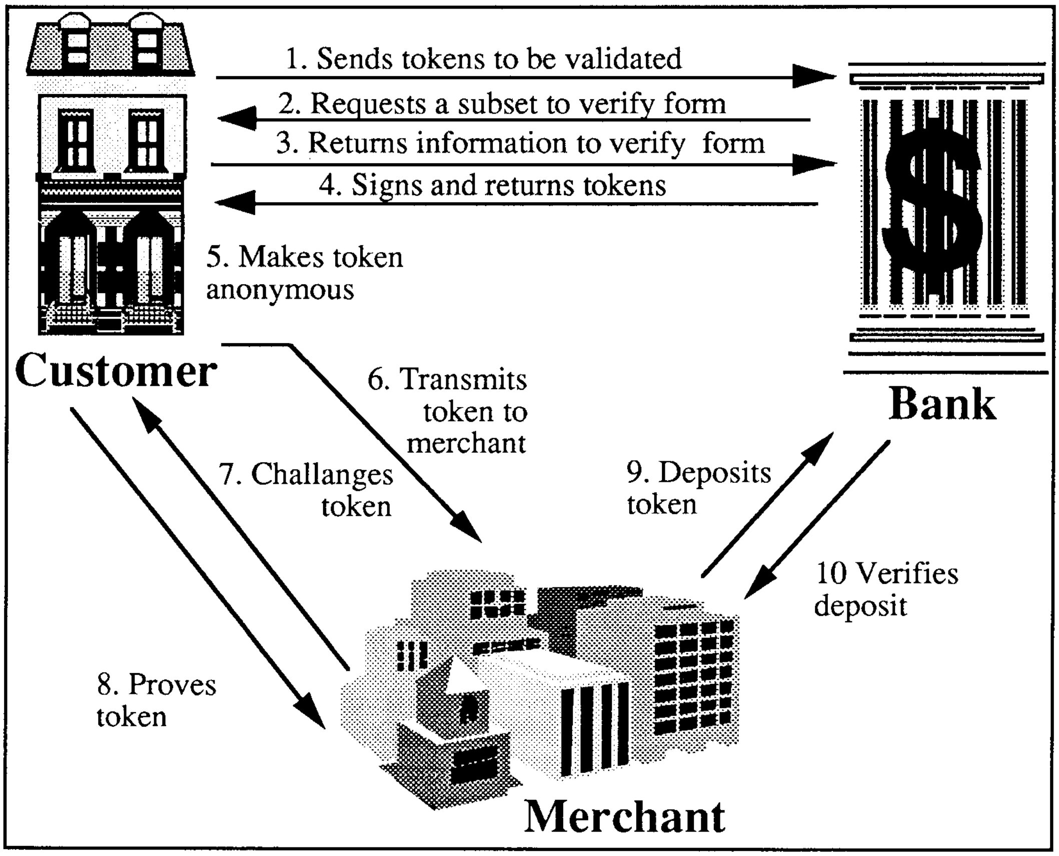

1. Sends tokens to be validated
2. Requests a subset to verify form
3. Returns information to verify form
4. Signs and returns tokens
5. Makes token anonymous
Customer
6. Transmits token to merchant
Bank
7. Challanges token
9. Deposits token
10 Verifies deposit
8. Proves token
Merchant

Figure 3: A Digicash Transaction

In the first step the customer formats a series of potential tokens. This requires two well known one way functions, g and f. (All operations are modulo n, where the bank knows the factorization of n.) The form of the token is

r_i^3, $f(x_i, y_i)$

where

$x_i = g(a_i, c_i)$ and $y_i = g(a_i$ XOR (account number || (counter + i), $d_i)$

Here || refers to concatenation and r_i is a random number. The account number and counter are known to both the bank and the customer.

After the customer sends a number of these tokens to the bank, and the bank selects a subset of these and returns them to the customer. This is shown in step two. In step three the customer must return the elements of the token to the bank (the appropriate r, a, c and d) so that the bank can assure that the other unread tokens presented are in the same form. Selection of a subset for verification is called *cut and choose*. Of course, there is a chance that the customer can obtain a signature on a problematic token which will enable the customer to commit fraud. The bank can determine its willingness to accept the risk of fraud, and implement the cut and choose technique in a manner consistent with that level of risk aversion.

The bank then signs the tokens that were not selected to be unmasked by the customer. The bank returns the signed token. The customer divides by the random number r_i, thereby obtaining a properly signed but anonymous token. The consumer and the bank then increase their corresponding counter values appropriately.

Next the customer selects an item which costs a number of tokens. For each token, the merchant can request either the appropriate a, c, and y values, or the appropriate x, (a XOR (account number || (counter + i)), and d values. Either set of these values allows the merchant to verify the token. However, with both of these the merchant could deconstruct the token and identify the customer through her account number. The merchant sends the token and the values used for verification to the bank.

Digicash is not money atomic. For every token that the customer spends there is a 50% chance that the next merchant will request a different set of values for token validation. If a different merchant asks for a different set of verification values, then the bank will have enough information to identify the account, and therefore the account holder. Thus, double spending is limited by the ability of the bank to detect the fraud, and the correspondingly high penalties.

3.1.2. Business Model

Digicash provides a mechanism for electronic payment. Digicash protocols do not provide mechanisms for discovery, negotiation, delivery or conflict resolution. The scope of Digicash is both its strength and weakness. The advantage is that Digicash can provide an elegant and simple protocol. The disadvantage is that Digicash cannot offer to decrease the cost associated with collection and dispute resolution. In fact, Digicash is specifically designed to mimic cash so that only the purchase itself and the detection of counterfeits are properly the business of Digicash. Thus any cost of fraud is transferred to the customer. This business assumption may not be valid, especially in the United States due to the Electronic Funds Transfer Act. (The Electronic Funds Transfer Acts specifically limits consumer loss in electronic funds transfers to $50 per lost instrument. It is not certain if a Digicash account meets the definition of an instrument.)

In Digicash, it is always assumed that the customer is the dishonest party. Since most credit card fraud results from unauthorized use of cards through theft or loss where the owner of the account cannot prevent the fraud, the validity of this assumption is questionable (Ballard, 1994). If the customer is indeed committing the fraud, the assumption that detection after payment is sufficient to reduce risk is also questionable, given the opportunity to recycle funds and disappear after successful fraud occurs (McClellan, 1995).

Recall the distribution of costs in a credit card transaction (Figure 2). In Digicash, customer billing is not possible since the customers cannot be identified. As with all Internet systems, card issuing does not exist and therefore cannot create cost. Customer authorization is unnecessary since each token is self-authorizing. Over limit and collections issues do not apply, since Digicash is debit only. Servicing, promotion, administration and processing are all completely automated so expecting cost decreases is reasonable. This leaves the cost of fraud as the potentially dominant issue.

3.1.3. Security

Digicash transactions have high privacy and low transactions cost. However, Digicash transactions may be subject to a high fraud rate.

Digicash fails to fully address merchant fraud. If a merchant receives a token and then deposits it, the merchant can claim not to have received the token. However, in this case the customer can provide the corresponding a, c and d values and thus illustrate to the bank that it is indeed the customer's account number embedded in the token. This means that a customer can prove payment at the cost of loss of privacy. However, Digicash creates no record of any sales agreement or delivery between the merchant and the customer. This means that the customer cannot prove that fraud occurred. In fact, if the merchant claims to have lost the token and the customer spends it again, the customer is at risk for fraud prosecution. Similarly, if a customer loses Digicash tokens, it is unlikely that the thief cares if the account owner is identified as a result of the thief's double spending.

The loss of the security of a Digicash server is very unlikely. However, if a bank's secure server was undermined, the attacker could generate an indeterminate number of valid tokens. Digicash is uniquely vulnerable since the account number in the generated token need not be valid. Therefore there is no assurance that a customer would eventually discover the loss, as is the case with notational currency systems. Thus there is the risk of long term undetected subversion of a Digicash server.

3.2. NetBill

3.2.1. A NetBill Transaction

A NetBill transaction can cover all phases of a purchase. NetBill includes secure price negotiation, final selection, payment, delivery and customer support. NetBill is optimized for purchase of electronic information goods over a network. A NetBill transaction is shown in Figure 4 below (Cox, 1995).

NetBill uses both public and private key encryption. NetBill uses the Digital Signature Standard (National Institute of Standards and Technology, 1991), RSA (Rivest, 1987) and Kerberos (Jennifer, 1988; Miller, 1987). NetBill uses public keys to generate and share symmetric keys to reduce the number of computationally expensive public key operations.

Before the NetBill transaction begins, a customer obtains a modified Kerberos ticket and symmetric key from the merchant. The modified Kerberos ticket is the equivalent of identity verification for the purposes of the merchant and the customer. It is referred to as the customer identity in later messages sent in the transaction to the merchant. The first message sent by the customer includes customer identification information necessary for a Kerberos ticket, addressee and an initial shared key. The customer encrypts the message in the merchant's published key and her own private key. This double signature assures that any response containing enclosed information must come from the merchant.

The merchant replies with a message containing a shared key and a Kerberos ticket. These are encrypted with the key initially sent by the customer. Now that a symmetric key has been established, the transaction itself can begin.

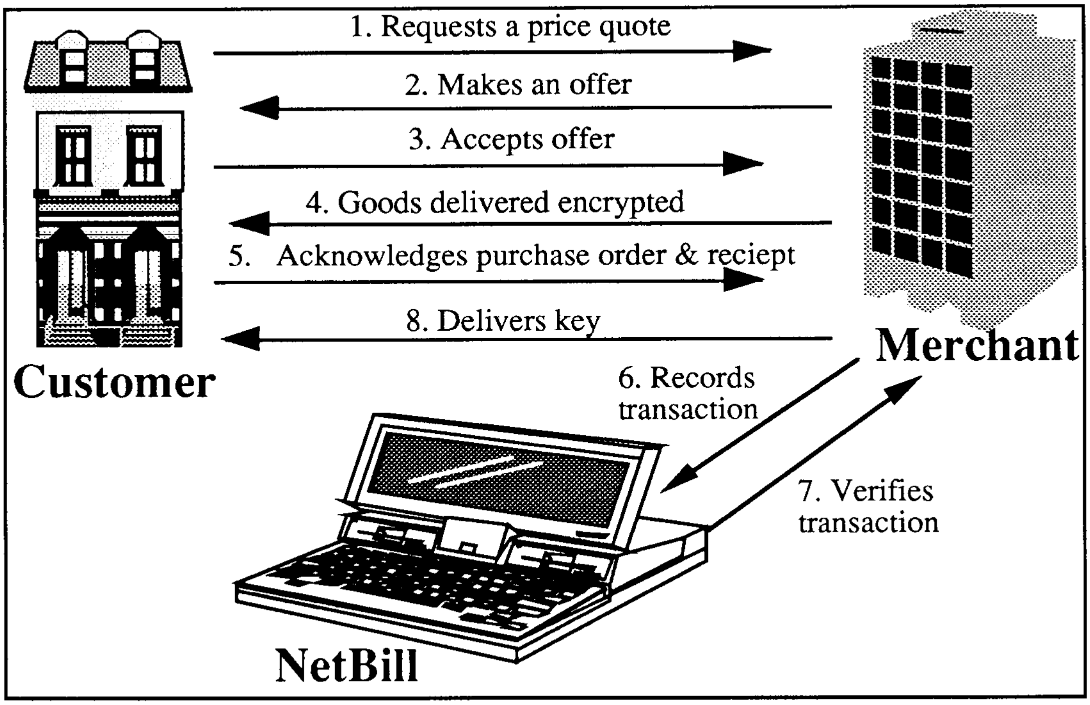

Figure 4: A NetBill Transaction

The first message shown in Figure 4 assumes the existence of the symmetric key. This message includes the customer's identity (the Kerberos ticket), a price offer, information about the requested item or purchase, and a transaction identifier. It may also include electronic coupons and membership certifications for appropriate discount or subscription verification. The merchant responds, in the second message shown above, with a product description, offered price and a transaction identifier. Both message formats are extensible through the use of request flags. These offer and request for information steps may be repeated several times.

If the customer decides the price is acceptable, she then requests the item with a message that includes her identity and the transaction identifier. This is step three. Again this message is encrypted with the shared key. The merchant then sends the goods, (step four) which are encrypted with a new key which is used only to encrypt these goods. The goods are accompanied by a checksum so the customer can prove that the merchant indeed sent exactly these goods. The merchant also sends an electronic invoice which includes the merchant's identity, a time stamp and a serial number. This serial number is globally unique and is used by the NetBill server to index transactions.

The fifth message is the customer's electronic purchase order. The electronic purchase order includes customer identity, the product description, negotiated price, merchant identity, the checksum of the goods, the checksum of the original request (step one), the checksum of the customer's account number, an account verification nonce and the electronic invoice from the merchant. In addition, the customer encrypts NetBill authorization information, any coupons or credentials used, account number, a nonce and a personal memo field. The customer signs the entire purchase order and sends it to the merchant. Note that the NetBill authorization information is unreadable by the merchant.

The merchant receives the electronic purchase order, adds the key to the goods, his account number, his identity from a NetBill Kerberos ticket, possibly his own memo field and then signs this new message. NetBill then extracts the payment information and sends a verification of payment, as shown in step six. Finally the merchant sends the key for the goods to the customer in the final step.

For any account debit the customer can demand from NetBill the signed electronic purchase order. This implies that NetBill has a contractual responsibility to provide refunds if NetBill cannot prove the account transfer was properly authorized. This also suggests that the customer and merchant need not trust NetBill unconditionally, since NetBill cannot make irrevocable commitment on their behalf without their signatures. The inclusion of nonces and time stamps in the messages above prevents simple replay attacks. NetBill is goods atomic.

3.2.2. Business Models

NetBill is targeted at a specific market: purchase of information goods on-line. The market for on-line information goods is hampered by the fact that goods are widely distributed and often have very low value. Many on-line information merchants are not large enough to have merchant accounts with credit card companies. Besides the number and small size of many information providers, this is a problematic market for current Internet commerce protocols because value of these merchants' items is so low, consumption happens soon after delivery, there is no standard for proof of delivery on-line, and there is no physical presence. NetBill is designed to reduce transactions costs by using the factors which make network goods difficult to purchase. NetBill is designed so that any consumer with a bank account can be an information provider using NetBill merchant software.

NetBill is designed for low price goods. NetBill also has a non-certified delivery technique for zero price or free goods. This is used for zero priced goods, like additional issues after the purchase of a subscription, or targeted coupons. This enables merchants to distribute these goods without being concerned that they will be made available to observers during transmission.

NetBill provides aggregation services as an intermediary. Aggregating transactions of ten and twenty cents into ten and twenty dollar charges results in orders of magnitude of cost spreading. Since NetBill is an intermediary, the marginal cost of credit acquisition for NetBill will be negligible. Currently NetBill is a debit system. Therefore over limit and collections are not issues.

NetBill automates authorization, customer service and many cases of fraud claims. Again card issuing is not an issue. Promotion is also automated, since NetBill is aimed at the online consumer.

NetBill reduces the cost of account acquisition and credit processing by accepting standard methods of payments through banks. NetBill provides per-transaction authorization and transaction aggregation using customer credit card or bank accounts.

The business plan of NetBill also makes clear that the provision of clients, servers and transaction processing should be subject to competition. By using open standards, NetBill can prevent any one server or software provider from becoming a bottleneck. Similar considerations drove Mastercard to create a system based on open standards (Mastercard, 1995).

Because NetBill can provide verified orders, this protocol could be used to provide verifiable receipts for orders of physical goods over the Internet. This would require the use of current verified delivery techniques, such as registered mail, for physical delivery. Currently, the extension of NetBill for verification of purchase orders for physical goods is not under consideration. Clearly NetBill cannot provide goods atomicity for physical goods.

3.2.3. Security

If a NetBill transaction is interrupted before the fourth step there has been no purchase. The customer has goods, but they are unreadable and therefore worthless. The customer can begin the transaction again. If the NetBill protocol is interrupted at step five the customer would have to again send the payment information. Recall the payment information includes the globally unique serial number. Therefore if the merchant forces the customer to send a second copy of the payment information the customer will not be charged twice. If the transaction is interrupted at step six the merchant is never paid and the customer never receives the merchandise. If the transaction is interrupted at step seven both the merchant and the customer can poll NetBill. The merchant can confirm payment. The customer can confirm payment and get the key to decrypt the goods. Forcing a failure at any step in the process does not allow any party to defraud the other.

The customer cannot verify the actual item delivered until step eight. The content of the item delivered at step four cannot be determined. However, the customer has a contract signed by the merchant in the purchase order. NetBill resolves disputes using this documentation. Recall that NetBill receives double-signed checksums that can verify the item delivered, and copies of the original request that verify the item requested.

Disputes over quality of merchandise are inevitable, particularly in information. Complaints may be as vague as issues of taste or as specific as failure of software to perform. This is particularly a problem with electronic purchases. In an electronic purchase the customer cannot view the item beforehand. To limit merchant fraud, NetBill tracks complaints against merchants to assure that merchants are not misleading customers.

If a NetBill server is subverted, the NetBill attacker could have up to one month to change accounts and abscond with funds. This is because it could take one account-activity reporting cycle for the first customer to complain of unauthorized debits. If merchants are not credited until customers approve transactions, then loss of server security would be costless. The financial security of the NetBill server depends on funds availability policies which are yet undetermined. In any case, NetBill keeps sufficiently detailed information for recovery from a fraud so that liability could be reliably assigned.

NetBill customers have very little privacy from NetBill, but can purchase their privacy from merchants. A customer can choose to purchase the service of a pseudonym provider. With the use of credentials, consumers may remain pseudonymous and still obtain any earned discounts. Regardless, the only information not available to the NetBill server is the item(s) purchased. NetBill knows the parties, date and amount of all transactions. Neither NetBill merchants nor NetBill servers are prohibited from compiling and selling customer information.

3.3 First Virtual

3.3.1. A First Virtual Transaction

A First Virtual transaction is shown in Figure 5 below. First Virtual was designed for the delivery of low priced information goods over a network.

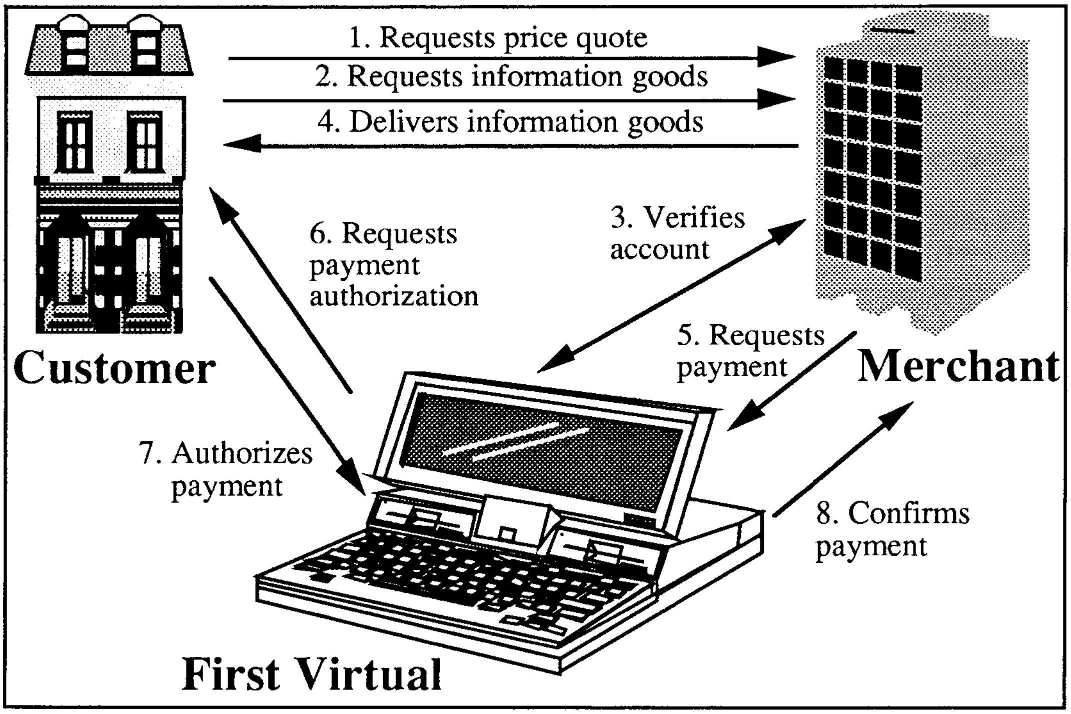

Figure 5: A First Virtual Transaction

To obtain a First Virtual account a customer first sends email to First Virtual that includes a customer-selected password. The customer then calls First Virtual and provides credit card information over the telephone. The credit card information itself is never sent over the Internet. The password is used by customers to access their First Virtual accounts.

After establishing an account with First Virtual, customers can begin making purchases. The customer selects an item from the Web page of the merchant. When the customer then requests that item the customer includes the First Virtual account identifier. The merchant contacts First Virtual and verifies the account identifier and password provided by the customer. The merchant is then contractually required to send the merchandise. After the merchant sends the merchandise the merchant sends the customer's payment authorization to First Virtual and requests payment, as shown in step five. First Virtual then sends an email message to the customer for authorization of the charge.

The email in step six and the request in step two travel through different parts of the Internet, like a telephone call to Tokyo and a fax to New York. Therefore First Virtual considers these independent channels. While it is simple to obtain a packet containing ordering information from First Virtual, intercepting the authorization request message to the customer is difficult. It would require either filtering every message received by the customer or breaking into the customer's home email account. Furthermore, there is no gain in completing the second, more difficult, part of the process because any attacker has already obtained the goods in step four. So it is likely that the email sent to the customer results in a valid reply in step seven.

First Virtual has money atomicity but not goods atomicity.

3.3.2. Business Models

The business model of First Virtual is based on three fundamental assumptions
- no credit card numbers are ever on the Internet,
- no replay attacks are possible, and
- the losses of a merchant who is unpaid for network-delivered information goods is negligible.

Fist Virtual is a protocol for the first generation of Internet commerce. As with all on-line systems, First Virtual has automated customer support, promotion, administration and processing. First Virtual transactions are large enough that aggregation is unnecessary. The goal of First Virtual is not to decrease the cost of a transaction by an order of magnitude but rather to provide immediate access to customers on the Internet for medium-priced information goods.

Insecure commerce has limited application. The size of a purchase with First Virtual is limited by the merchant's tolerance for fraud. Merchants with high quality goods for which there is a high demand are unlikely to accept high levels of fraud. First Virtual works well for low priced goods with a small to medium market, or high priced goods with a specialized market.

First Virtual is only for information goods delivered over the Internet. Because no First Virtual order is binding it cannot be used to make verifiable orders for physical goods. First Virtual's approach allows every Internet user to be both a merchant and a consumer. This vastly expands the number of possible merchants and therefore the probability that there will be information of interest to a customer.

3.3.3. Security

First Virtual is very clear that the credit card information itself is never sent over the Internet. Thus the very lack of widespread interoperability between forms of network commerce is an advantage for First Virtual, since you cannot trade First Virtual account authorization for any other financial instrument.

This is not to suggest that First Virtual is secure. An attacker need only trap a packet which has the name of a First Virtual account holder to receive information free. Since there are well-known locations which receive many of these packets (for example, the First Virtual Infohaus), finding such a packet is unlikely to be difficult.

Merchants can get customer identity information but not customer credit card information. Merchants do get the information necessary to authorize further purchases. However, merchants will not profit, so this crime is unlikely.

Customers have no privacy in First Virtual. First Virtual gets complete information about a customer's purchasing habits. Customers cannot make anonymous purchases. There is no restriction on First Virtual's or merchants' compiling and selling customer data. First Virtual merchants are contractually required to keep detailed transaction records for at least three years after the transaction (First Virtual, 1995). Since all messages are sent in the clear a curious attacker could develop a profile of a customer, or of a particular merchant's transactions.

3.4 Secure Electronic Payment Protocol

3.4.1. A Secure Electronic Payment Protocol Transaction

The Secure Electronic Payment Protocol (Mastercard, 1995) provides various levels of protocols for those with different Internet access capacities. To be consistent with previous protocols, I will discuss a transaction for a customer with Web access. The Secure Electronic Payment Protocol is designed for the purchase of any medium or high priced physical or electronic good over the Internet.

A Secure Electronic Payment Protocol (SEPP) transaction is shown in Figure 6 below. The bank shown below is the merchant's acquirer.

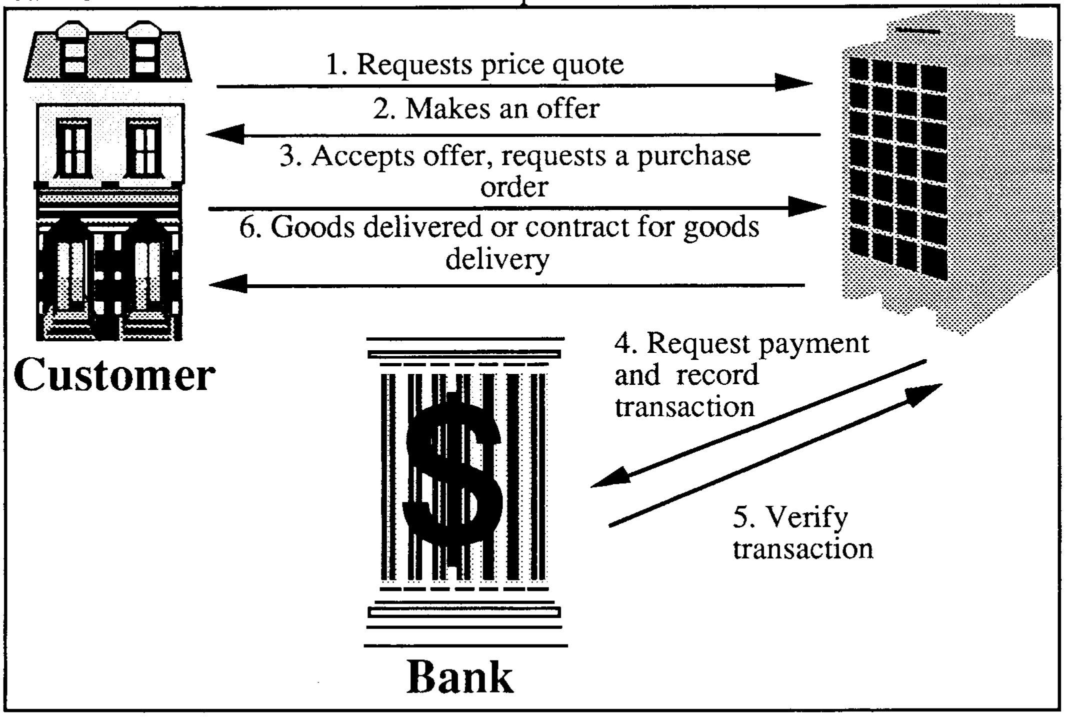
1. Requests price quote
2. Makes an offer
3. Accepts offer, requests a purchase order
6. Goods delivered or contract for goods delivery
Customer
4. Request payment and record transaction
5. Verify transaction
Bank

Figure 6: A Secure Electronic Payment Protocol Transaction

Secure Electronic Payment Protocol uses RSA (Rivest, 1987) encryption and X.509 certificates.

Before a transaction begins merchants and customers must obtain verification of their right to use the SEPP protocol. This is done by obtaining certificates from a trusted authority, called the certificate management system by SEPP, which includes their respective publicized keys. These keys and certificates verify the identity of the various parties and provide nonrepudiation.

The SEPP transaction begins with a request for a price quote (called an Initiate message). The Initiate message includes a cardholder identity (the certificate) and a transaction identifier. The transaction identifier does not need to be globally unique; it is used only by the customer to associate messages with a transaction. The merchant responds with an invoice which contains the customer's transaction identifier, the merchant's transaction identifier, and data requested by the initiation, such as price, items and specifications.

A customer notifies a merchant of the intent to make a purchase using a Purchase Order Request message, as shown in step three. The customer generates a purchase order using the previously received offer. This new message includes a hash of the purchase order signed by the customer, purchase information encrypted in the published key of the bank, cardholder identification, the merchant transaction identifier and the customer transaction identifier.

In step four the merchant requests payment by forwarding the customer's encrypted purchase information to the bank. In addition to the purchase information, the merchant includes the merchant's and customer's identity certificates, as well as customer and merchant transaction identifiers. The merchant also includes a hash of the purchase order signed with his private key. The signed purchase orders from the merchant and the customer provide nonrepudiation.

The bank authorizes the merchant's identity, and then requests payment from the cardholder's bank. The bank then sends a signed verification to the merchant.

SEPP has money atomicity, but not goods atomicity.

3.4.2. Business Models

NetBill and First Virtual are financial intermediaries that provide preprocessing for off-line acquirers. Arguably, NetBill and First Virtual are merchants from the perspective of the acquirer. With both Mastercard's Secure Electronic Payment Protocol and the VISA Secure Transaction Technology (VISA, 1995) protocol, the acquirer would be on the Internet. This is feasible for the obvious reason that there is no need to aggregate large charge card purchases made over the Internet. The same customer support, order processing, administration and promotion savings can be obtained by Mastercard and VISA. The Mastercard and VISA protocols may not compete as much as complement the approaches of the previously mentioned Internet commerce providers.

Mastercard models Internet commerce as mail order and telephone commerce. (This is the obvious implication of the fact that the merchant takes the risk for invalid purchases, as in mail and telephone orders, rather than acquirer, as is the case with purchases with physical presence.) The Mastercard protocol differs fundamentally from other protocols in that only traditional merchants are allowed to sell goods. This means that small publishers and professionals working at home cannot use SEPP if they do not have merchant Mastercard accounts.

Mastercard has chosen to develop an Internet protocol using the traditional open Internet process of issuing drafts and requesting comments. SEPP will be built upon standard Internet protocols. In contrast, the Secure Transaction Technology attempts to leverage the dominance of Microsoft operating system to popularize its technology.

3.4.3. Security

The most dramatic improvement of the Internet protocol over the mail order and telephone protocol for Mastercard is that the merchant gets enough information for only one purchase. Unethical merchants cannot use SEPP information for replay attacks. If the transaction identifiers are monotonically increasing, for example by being a function of time, then there is no possibility that the same transaction identifiers would be repeated. Even if the transaction identifier were generated in a deterministic way, the merchant could not produce a purchase order signed with the customer's private key.

The SEPP protocol does not include negotiation or verification of delivery of information goods. A customer can claim not to have received goods already consumed, and a merchant could claim to have provided goods not sent. Therefore the security of SEPP depends upon the delivery mechanism used. The strength of nonrepudiation is limited when fulfillment of that promise cannot be confirmed.

The Mastercard protocol provides more privacy than current credit card transactions, since the customer's financial information is hidden from the merchant. The Mastercard is less private than NetBill since Mastercard knows the item(s) purchased. It is more private than First Virtual since the merchant is not apparently required to maintain records of customer purchase for three years. This is made unnecessary by the nonrepudiation enabled with public key cryptography.

4. Conclusions

Different Internet commerce systems offer different trade-offs. Different commerce protocols are suited for different applications. Table 4 below shows the commercial characteristics of different protocols.

Internet commerce Protocol	Transaction Cost	Customer Tolerance for Fraud	Merchant Tolerance for Fraud	Customer Privacy
Digicash	low	high	low	high
NetBill	low	low	low	medium
First Virtual	medium	low	high	low
SEEP	medium	low	medium	low

Table 4: Commercial Characteristics of Internet Commerce Protocols

In comparison a current credit card telephone order has high transactions cost, requires that both merchants and customers have a high tolerance for fraud, and provides the consumer low privacy.

Information disclosure differs in the Internet commerce systems as well. Consumers in the United States are increasingly concerned about privacy (Camp, 1995a). In Canada, under the Freedom of Information and Protection of Privacy Act, the Privacy Commissioner is required to determine if a business practice violates privacy and to act to end any privacy violations after an inquiry. In the European Community consumer information is protected by the EC Directive 95[1]. Together Canada, nations of the European Community and the United States contain eighty one percent of the hosts on the Internet (calculated from data obtained from IDS, 1995b).

[1] Directive 95 of the European Parliament and of the Council of the European Community on the protection of individuals with regard to the processing of personal data and on the free movement of such data was approved July 20, 1995.

Internet Commerce Protocol	Information Available to Observers	Information Available to Merchants	Information Available to Transaction Processor
Digicash	Merchant, amount[2], time, item[2]	Amount, time, item	Merchant, amount, time
NetBill	Time	Customer[3], amount, item, time	Customer, merchant, amount, time
First Virtual	Merchant, customer, amount, time, item	Customer, amount, time, item	Customer, merchant, amount, time, item
SEPP	Merchant, customer, amount, time, item	Customer, amount, time, item	Customer, merchant, amount, time, item

Table 5: Information Availability in Internet Commerce Protocols

Clearly these protocols are suited for different merchants, products and consumers. The decision to select a particular protocol is a function of customer sensitivity to data surveillance, merchant desire for information, merchant sensitivity to competitors' data surveillance, and customer and merchant tolerance for fraud.

The two tables above illustrate that decreased privacy alone does not yield increased security. In fact, for those transactions where the content of the transaction has value to possible hostile observers, decreased privacy decreases security. The previous analysis further illustrates that there is a wide range of variables which are a function of consumer and merchant preference. Of course these preferences may change over time, as defrauded parties or parties subject to surveillance become more sensitive to issue of security and information availability.

[2] This is a function of the negotiation protocol. Customers and merchants could choose to encrypt purchase information.

[3] Recall the customer can purchase psuedonymity.

5. Bibliography

Camp, L. J. ,Sirbu M. & Tygar, J. D., 1995a, “Token and notational money in electronic commerce", *Usenix Workshop on Electronic Commerce*, July, New York, NY, proceedings in preparation.

Camp, 1995b, “Options, Opportunities and Obstacles in Electronic Commerce”, The Conference on the Future of Electronic Banking, Columbia University, New York, NY, October 23.

Cerf, V., 1993, “How the Internet Came to Be”, *The Online User's Encyclopedia*, ed. B. Aboba, Addison-Wesley, New York, NY

Chaum, D., 1985, "Security Without Identification: Transaction Systems to Make Big Brother Obsolete", *Communications of the ACM*, No 10, Vol. 28, pp. 1030-1044, October.

Chaum, D. 1988, Untraceable Electronic Cash, Advances In Cryptology - Proceedings of Crypto ‘88, ed. S. Goldwasser, Spring-Verlag, Berlin, pp. 320- 327

Cox, B., Tygar, J.D. & Marvin, S., 1995, “NetBill Security and Transaction Protocol”, Usenix *Workshop on Electronic Commerce*, July, New York, NY, proceedings in preparation.

Financial Service Technology Consortium, 1995, *Electronic Payments Infrastructure: Design Considerations*, http://www.llnl.gov/fstc/projects/commerce/public/epaydes.htm, November

Jennifer G., Steiner, B., Neuman C., & Schiller.J.I., 1988, “Kerberos: An Authentication Service for Open Network Systems”, *Proceedings of the USENIX Winter Conference,* February, pp. 191-202

Haber, S. & Stornetta, W.S., 1991, “How to time-stamp a digital document.”, J*ournal of Cryptology*, Vol. 3, pp. 99-111, 1991.

Hanushevsky, A., 1995, *Electronic Commerce Page*, http://abh.cit.cornell.edu/ecom.html, November

Internet Domain Survey (IDS), 1995a, *Exponential Growth of the Number of Computers Connected to the Internet*, http://www.nw.com/zone/WWW/top.html, November

Internet Domain Survey (IDS), 1995b, *Hosts Stats by County,* http://www.nw.com/zone/WWW/isoc-pr-9501.txt, November

Mastercard, 1995, *Secure Electronic Payment Protocol Specification Draft Version 1.1*, http://www.mastercard.com/Sepp/sepptoc.htm, November , Part 2, pp. 10

McClellan, D.,1995, "Desktop Counterfeiting", *Technology Review,* http://web.mit.edu/afs/athena/org/techreview/www/articles/feb95/mcclellan.html, February/March.

Miller, B.C., Neuman, C., Schiller, J.I. & Saltzer, J.H., 1987, "Section E.2.1: Kerberos Authentication and Authorization System, " *MIT Project Athena*, Massachusetts Institute of Technology, Cambridge, MA., December

National Center for Supercomputing Applications, 1995, *NCSA Mosaic Web Index,* http://www.ncsa.uiuc.edu/SDG/Software/Mosaic/Docs/web-index.html, November 1995

National Institute of Standards and Technology, 1991, *Proposed Federal Information Processing Standard for Digital Signatures,* Federal Register, Vol. 56, August, pp. 42980-42982

Schneier, B., 1995, *Applied Cryptography, Second Edition*, John Wiley & Sons, Inc., New York, NY

Schwartz, M., 1987, *Telecommunications Networks Protocols, Modeling and Analysis,* Addison-Wesley, Reading, MA

Sirbu, M., & Tygar, J. D., 1995, "NetBill: an Internet commerce system optimized for network delivered services", *IEEE ComCon,* San Francisco, CA; March 6

Verisign, 1995, *Verisign Expands Digital ID Offerings To Leading Web Servers*, http://www.verisign.com/pr/pr_servers.html, November

VISA, 1995, *Secure Transaction Technology Specifications Version 1.1*, http://www.visa.com/visa-stt/index.html, November

Yee, B., 1994 *Using Secure Co-processors*, Ph.D. dissertation, Carnegie Mellon University. Available as CMU technical report CMU-CS-94-149

6. Glossary

clear sign or sign: a message accompanied by a hash value which can be used to verify that the message has not been altered and has been sent by the originating party, i.e. the hash value has been encrypted with the originating party's private key of his or her public key set

goods atomic: the transfer of funds and the transfer of payment are intrinsically linked; either both happen completely or neither happens at all

hash or checksum: a compressed form of a document which is constructed so that no information about the contents of the document can be determined from the has, yet the hash value of a document is unique for every document

heterogeneous: a network is heterogeneous if there is no standard software or operating systems; ex. a network with Unix workstations, IBM compatibles and Macintoshes is heterogeneous

money atomic: the transfer of money either occurs completely or fails completely; it cannot happen partially

nonhierarchical: a network is nonhierarchical if there is no recognized or central authority to impose standards or prioritize machines or tasks

nonrepudiation: an action can to be shown to have been taken by an individual; ex. a customer provides nonrepudiation to a bank by signing a traveler's check

public key or public key set: a set of keys which together enable using and inverting a specific function for the purpose of encryption, i.e. information encrypted with the secret key can be decrypted with the matching published key and information encrypted with the published key can only be decrypted with the secret key.

published key: the set of numbers made public for an individual for use with a public key algorithm private key or secret key: the set of numbers kept secret for use with a public key algorithm; the security of a public key algorithm depends on the secrecy of these keys

replay attack: the transmission of previously obtained information by an unauthorized individual, this is the attack used when a consumer's credit card number is used after disclosure without authorization

shared key or symmetric key: a number or set of numbers used in a symmetric encryption algorithm; i.e. an algorithm where the same key encrypts and decrypts information

subscriber: a person that receives messages sent out to a dlist in his or her personal mailbox

telnet: connection to a physically distant machine, requires terminal emulation on the user's machine