Editor: Nevenka Dimitrova
Philips Research

Bandwidth as Currency

Scott Moskowitz
Blue Spike

While cash and its equivalents let us objectify commercial transactions, *money*, a medium of exchange, is simply information.[1] Provisioning bandwidth in an actuarially consistent manner has many far-reaching implications reminiscent of the commercial rationalization of railways in the 1870s—in that period's *path to profitability*. Mapping packet flow to willingness to pay, as rail cargo content was matched to the cost of carry, will better utilize bandwidth, enhancing network value consistent with the information being exchanged. The notion of a *packet watermark* can enable bandwidth as currency. Payment facilities are generally treated as separate, independent data flows from the actual data being transacted. Packet watermarks map payment facilities to the fidelity, discreteness, or functionality of the data demanded, representing a consistent means of determining a willingness to pay. This mapping acts as a receipt for data commerce. Uniquely identifying the exchange of objects representing *abstractions of value*[2] enhances a transparent, liquid information economy.[3,4]

The basic framework of areas that would enable this bandwidth provisioning includes

- efficient packet provisioning on a network using a packet watermark,
- unique identification of bandwidth availability and flow,
- bandwidth credentials creation to enhance liquidity and derivative pricing for future estimated use of bandwidth, and
- cryptographic protocol-based rules that let market mechanisms objectively bill and subsequently resolve disputes.[5]

Current information commerce of media or functionally rich data objects typically lacks any assessment of responsibility for the parties and intermediaries handling data objects. Applying cryptographic uniqueness to the packets and traditional cryptographic key-based watermarking to the data objects would uniquely identify the object and monitor the information exchange.[6]

Packet watermarks differ from traditional digital watermarks. Whereas digital watermarks act on the application-layer data object intrinsically related to type and use of the data,[7,8] packet watermarks relate to the actual transmission. Packet watermarks assist with the authenticated provisioning of packet flows between users, can break the actual transmission into parts, resequence the parts, and introduce additional communication-related information—which can later be associated with each other. Preferably, the packet watermarked data won't interfere with the traditional digital watermarks, which establish responsibility for the objects being transacted.

As with other transmissions, end users don't care about the nature of the packets. However, they can benefit from using the best paths for getting information. Vendors offering information could use packet-watermarking applications to objectively assess responsibility for data—for legal or economic reasons. This could also avoid double payments of bandwidth, where vendors handle the sending and receiving costs, instead of an optimized path between a sender and a receiver. These applications could enhance trust in entities (that is, devices and people) that are increasingly associated with some intangible, yet recognizable information associated with the transaction itself. *Trusted computing* might result from these approaches.

Tragedy of the network commons

Bandwidth suffers from the tragedy of the commons, that is, everyone wants to increase their quota, even if it's detrimental to the global system. We need transparent, liquid, and secure

1070-986X/03/$17.00 © 2003 IEEE

Published by the IEEE Computer Society

protocols to accurately assess and provision bandwidth in a manner consistent with responsible information commerce. Some refer to this issue from a public policy perspective as *spectrum management*, others as *the future of ideas*.[9]

Addressing the optimized allocation of bandwidth has largely been the domain of quality of service (QoS) approaches, competitively offered in tandem with traditional peering arrangements between large carriers. Early work resulted in caching technologies, which push higher demand data closer to the access points for which the data is demanded. QoS attempts to make decisions about bandwidth accessibility based on a user's ability to access information within some predetermined time frame. For instance, if X number of users can access Y amount of bandwidth over some fixed period of time T, we can estimate bandwidth as a function of satisfying users X, or some percentage of X, for each increment of Y divided by T.

Traditional telephone billing systems provide a somewhat accurate measure of bandwidth use, measured as discrete instants of time, and the general, or hybridized, path by which users are connected. However, present information commerce of media or functionally rich data objects typically lack any assessment of responsibility for the parties and intermediaries handling these objects.[10] Blue Spike has developed a number of novel concepts based on 10 years of research and development in intellectual property rights management and cryptographic payment systems.

While priority of transmission paths helps alleviate bottlenecks within a given network, mapping demand for bandwidth has become increasingly difficult. This might result from user's assigning a high priority to their data. It also could be a result of competing interests within the Internet service provider space.

Several technological approaches to the bottleneck issues attempt to minimize computational overhead. Data compression schemes for media-rich content that support streaming or sharing an audio or video signal—for example MPEG-4—reduce the total number of bits transmitted over the communication channel. In the functional data space, optimized languages attempt to reduce computational overhead for a variety of applications, including multimedia messaging services or Java midlets. The reality is, not all data in a particular format or market segment carries equal commercial value. The market's participants have a deficit of time to offer and enter into commercial transactions, placing a premium on accessibility and satisfaction of good or service demands close to real time. Networks should let market mechanisms assist in providing and pricing data that's consistent with the bandwidth requirements and the rights of content or software creators while not interfering with the consumer's experience in transacting data.[11]

Internet protocol (IP) provides each networked device with an IP address. IP version 4 (IPv4) incorporates option fields that can be exploited at any place in the transmission chain for writing/embedding and detecting/recovering a specialized type of digital watermark that's suited for provisioning and pricing schemes, bandwidth prioritization, management systems, and dispute resolution and clearinghouse functions. Because of the sequential nature of TCP/IP, network researchers have suggested assigning higher priority to the perceptibly significant data in a data object.

Nonsequential transport for bandwidth provisioning

One way to optimize data transmission speed is based on Reed–Solomon error-correction coding. TCP/IP packets represent predetermined packets of data, that is, they have a specific size without regard to the data object being rendered. Therefore, coarser estimates of the data objects' aesthetics or signal characteristics let mathematical values be assigned to a larger portion or subset of the data object. A simple linear equation can define the independently derived values representing the data object. These mathematical values represent groupings of packets that aren't sequentially ordered but fitted to the characteristics of the data object being broken down for transmission. Additionally, systems or devices related to sending and receiving data can handle these values to speed data transmission.

Data chunks aren't sequential with error-correction coding, as it is with TCP, but are generated with variations on the Reed–Solomon code. As a result, receivers of the data get transmission chunks that can be reconstructed nonsequentially, but efficiently, so long as they receive assigned data values. The chunks may also overlap the packets that would typically represent the object. On the receiving end of the transmission, some applications first reconstruct those data signal features deemed perceptibly significant. Medical data, which might be time sensitive, can

benefit from this form of transmission. This approach speeds the routing of data over a network in a manner consistent with the perceptible value of the data, but it still lacks an effective way of attributing responsibility over data transmissions.

Tiered bandwidth quality with traditional digital watermarks

A wholly different approach combines traditional digital watermarks embedded in a full-bandwidth signal. These signals might have distortions or quality levels intentionally introduced that have differential pricing levels associated with predetermined keys for formulating a subset of the original signal's quality level and a rough estimate of overall signal quality demand (via the exchange of authentication information carried by embedded signals in the streamed data objects).[12] Each client would still receive a full-bandwidth signal at some level of quality up to full, and a yield in time measured via the verification of the embedded bits reported back to the server.[13]

Using transfer functions—which weigh the input to output of data—introduces degraded quality levels as a form of *chaffing* or *scrambling*. An approach that has a relationship with the signal's characteristics would not require separately handling and encrypting each quality level of a given signal served on a per request basis. Here, I discuss higher bandwidth granularity in observing the link between information, quality, and demand.

Business side of bandwidth provisioning

IPv6 includes proposals for additional optimizations. In contrast with current IPv4 systems optimized to handle end-to-end data transmission without regard for the data's content, IPv6 will enable traffic prioritization, low-level authentication with encryption, and better handling of audio and video streams. The labeling scheme discussed in this article enables better granularity in handling data packets with a labeling scheme over network infrastructures. The approach's authentication protocol prevents labeling fraud to reduce freeloading on paid bandwidth flows. The method uses packet flow watermarks differently than traditional digital watermarking. It prioritizes data traffic and defines the transmitted data so that it's consistent with the rights of the content or the data's functionality. The method also includes provisions for clearinghouse facilities and certification of traffic. Further, it offers secondary or derivative markets for assisting in efficient pricing of future bandwidth. From these novel techniques, I anticipate appropriate digital credentials for bandwidth pricing and use—called a *bandwidth credential* or *bandwidth digital certificate* as per traditional cryptological terminology.

We can now address market-based pricing of data in a manner that provides bandwidth efficiently. When a steganographic cipher or cryptographic-key based method watermarks a single data object, aesthetic or functional, it can be made unique.[14,15] Uniquely watermarking flows of packets, postage for packets (bandwidth provisioning) represents a natural extension for mapping granular commercial value of demanded packets versus other packets. By associating identifying and authenticating information of the watermark flows of packets, networks can more efficiently apportion bandwidth to meet market demands. The steps of identification, authentication, verification, and authorization are like negotiable levels of information exchange required by either party to a transaction. Certain types of transactions will require more or less information exchange than others, including higher security protocol demands to flexibly handle as many potential transactions as possible and bit commitments—as with zero knowledge signature schemes—by one of more of the parties for any additional assurance. More specifically, demand for information over networks and a better ability to identify the packets people are willing to pay for can be enabled in a highly efficient, cost-effective manner when demand is mapped to packets and their paths.

What also results is a better accounting system that provides billing packets to the appropriate parties and resolves disputes more objectively because cryptographic protocols assure a higher level of confidence in how provisioning is handled. Similarly, packet watermarking makes it possible to charge for bandwidth so that it resembles traditional telephone billing systems, albeit based on the value of data objects and the demands for the underlying packets in terms of time, quality, or functionality. The difference is that telephone billing systems don't consider the contents or paths of packets, nor do traditional telephone systems assist in creating a means for competitively evaluating bandwidth based on consumer demand for data. This demand can be compared to a more consistent media or in functional terms (type of media, associated rights,

authenticity of the data, quality level of the media based on a differential price, optimized functions, code or algorithms, and so on) and not solely on data size terms.

A network, thus enabled, can check and verify efficient bandwidth delivery on a packet level and can store information concerning better paths between senders and receivers. For certain economic or business models, further features can be added to make Internet handling of data similar to how billing works for traditional telecommunications companies. Such companies buy bandwidth resources in bulk and don't necessarily have any underlying understanding of what the bandwidth is used for, why it's being demanded, nor how to encourage higher value-added for any given bit for each bit per time calculation. The following describes one framework for measuring bandwidth:

- The intrinsic value $V_I = X$ x $(\min_0 - \min_1)$, is the money saved in telecommunications costs by using a higher bandwidth. The intrinsic value can be negative, implying a compensating premium placed on the time saved by using a more expensive transport. Note that $\min_0 \geq \min_1$.

- The percentage chance of failure represents the chance a user can't exercise rights (immediate purchase or sale of bandwidth) or option (where the option is the right, but not obligation to purchase the underlying asset) for bandwidth. If the probability of failure is P_f, where $0 \leq P_f \leq 1$, and the value of the right is V_0, in the absence of failure, then $V_f = (1 - P_f)V_0$.

- The convenience premium might apply to the particular or uniquely identifiable data objects, whether the data object is streamed, date or time schedules, geographic locations of either the provider or user, the hardware or software underlying the network, or some other unique circumstances including live performances. The more demand in excess of supply, the higher convenience C, will rise. V_C is then a function of supply and demand. Thus, $V_{\text{real}} = V_{\text{theoretical}} + V_C$.

- The time value is a function of the exercise period of a bandwidth right. It's proportional to P_f, since more time allows for transfer of recovery from an individual failure. There are two components of time: over what period a transfer can be initiated, and for how long the transfer can last once initiated. Thus, overall, $V = (1 - P_f)\ (V_I + V_T + V_C) = (1 - P_f)\ [(X(\min_0 - \min_1)) + V_T + V_C)]$ (Convenience premium V_C should be independent of all other values, except V.)

The pricing model also incorporates classic Black–Scholes options pricing, or derivations of this model, to price future value for bandwidth.[16] The following properties describe Black–Scholes: The standard deviation of the asset's value (in this case, bandwidth, or that which is optioned) multiplied by the square root of the time of the option's expiration. Essentially a ratio of the asset value to the present value of the option's strike price represents the underlying property of future price. The strike price is the price at which the option is offered and later exercised. To purchase or to sell is the difference in the right of the option and is called a *call* or a *put* (a put is the right, but not obligation to sell; a call is the right but not obligation to buy the underlying asset). More generally, the Black-Scholes equation is as follows:

$$C_0 = S_0\ N(d_1) - X\ e^{-rfT}\ N(d2)$$

Where

S_0 = the price of the underlying asset (a predetermined value)

$N(d_1)$ = the cumulative normal probability of unit normal variable d_1

$N(d_2)$ = the cumulative normal probability of unit normal variable d_2

X = the exercise price

T = the time to expiration or maturity of the option

rf = the risk free rate (a value that can be predetermined at the time of pricing the option)

e = the base of natural logarithms, constant = 2.7128…

$$d_1 = [(ln\ (S\ /\ X) + r_f\ T)\ /\ (S \div T)] + [1/(2\ S \div T)]$$

$$d_2 = d_1 - S \div T$$

Because the denominator (time) is fixed at any discrete moment, thus maximizing the economic value for the numerator (the bit) given a market for information goods and services, a higher economic value can be attributed to a given network that implement the features I describe here. While no one can know in advance the demand for a given data object, parties can agree to the

cost of bandwidth for a given business activity (such as streaming a live concert or handling bandwidth-based transactions tied to a subscription with a bandwidth device such as a cell phone). Streaming, to date, isn't economically viable because vendors haven't taken a packet level view of the flow of data to people demanding a stream. Nor have vendors tied payment or willingness to pay to the packets in a consistent manner with the data being consumed.

Ultimately, the notions presented in this article emphasize the different needs of providers and consumers of content and the multivalent nature of trust. So long as some preexisting payment or credit facility exists, decisions or policies regarding the level and detail of security or credentials should be made as flexible as possible regarding data and computational resources. Some transactions might only require a check sum not a more secure, independently verifiable cryptographic digital credential such as a digital signature with an ITU-T X.509 digital certificate. Combining verifiable identification inherent to digital certificates with bandwidth provisioning results in a bandwidth digital credential. For networked devices, payment facilities can easily be enabled and tightly integrated, especially if such devices have IP addresses or some similar uniquely attributable ID.

We can enhance tangible products with the unique information and transaction processing as a basis for serializing the actual article of manufacture. A major thesis of the techniques described here is that commerce must balance privacy with concerns about piracy. Further, commerce is about uniqueness or the receipts for copies sold, not originals, for which uniqueness may be nonverifiable. Recognition, not physical location, begets the commercial need for establishing responsibility over copies, whether aesthetic or functional data.

Packet watermarking

When a receiver requests a data object from a sender, the sender creates a packet flow with the receiver's address and sends it to the Internet. The packets might make many hops *in the cloud* before arriving at the receiver's IP address. At each node, a router examines the address and chooses a route to the next node. Often, there are many possible routes from each node to the final destination. These routes might be ranked by a number of criteria, including current load, historical load and reliability, and current and historical latency. All these factors could help route individual packets by more or less optimal paths—assuming that the router could discriminate between different flows. The packet watermark becomes the method by which the router identifies streams and creates differential QoS.

A packet watermark is cryptographically associated with the contents of the packet itself. An important issue is that the packets might contain functional data as opposed to aesthetic data. Mapping demand via cryptographic protocols to aesthetic data—in perceptibly significant portions of a signal—is only slightly different from mapping functionally significant data such as source, object, or executable code. For example, a traditional digital watermark might depend on the signal characteristics of the signal being watermarked. If watermarking occurs within a key-based system, a cryptographic association between the key and the signal or function via the watermark might exist. Besides the noise or signal characteristics in the signal, the key can be seeded by independent, random information to make it more difficult to decode, even if a potential pirate found the watermark in the signal.

Benefits from key-based watermarks are multifold. Key-based watermarks verify a data signal or object to establish responsibility for the signal or alert users of unauthorized copies. Similarly, a packet watermark sniffer could detect unauthorized use of a particular routing priority. The sniffer samples a fraction of the overall traffic to detect, and deter, abuse of the system. It reads the watermark on the packet, checks the authentication, and signals invalid packets. If necessary, the flow can then be rerouted or halted, depending on the terms of the commercial contract. Additional benefits can assist in a workable exchange that might further alert participants of particular users or unauthorized parties. This can prevent denial of service attacks and similar misuse of network traffic. Conversely, the exchange might maintain histories of the effectiveness of particular routes or particular parties that command a premium price or similar consideration for its recognition or reputation. For these reasons, an open form of rights and responsibility management—as opposed to traditional notions of access restriction-based digital rights management, are enabled for data—aesthetic, or functional content owned by other parties or for which rights need to be cleared.

Bandwidth provisioning

The packet watermark can help classify a data stream for a particular QoS. The data stream might be organized into a number of packets, and the sender can add a watermark to each packet's header. The size of the watermark can vary, but for illustration, a 32-bit watermark is stored, in say, the stream ID option field (that is, in the header) in the IPv4 packets. Preferably, the same 32-bit watermark would be placed in each and every packet in the flow. Additionally, for this example, the watermark's four most significant bits (MSBs) could help identify the QoS level, yielding 16 available levels, and the remaining 28 bits of the watermark could then uniquely identify the flow. One possible implementation for the remaining 28 bits is to store a unique identifier associated with a watermark packet key.

For example, the sender could create an array of the flow's secure hashes (for instance, SHA-1 or any hashing protocol deemed secure by the party or parties) using the watermark packet key. The watermark packet key, the watermark, and a portion of the flow make up the input to a hash function. The flow associated with one, two, or even more data packets could make up the portion of the flow used as input to the hash function. For this discussion, I consider the flow associated with one packet (that is, the portion of the flow inserted into one TCP/IPv4 packet). The hash's output might have a predetermined number of bytes. The array is the set of all hash outputs generated using successive portions of the flow until the complete flow has been processed. The outputs of the hash, the watermark packet key, and watermark are combined to create the watermark identification (WID).

Accordingly, the watermark can be matched to a corresponding WID. The component parts of the WID then help check the flow's authenticity. Moreover, if a portion of the watermark helps identify a particular QoS level, then we can evaluate the data for compliance for a particular path (such as for transmission by a compliant router). For higher security requirements, we can easily implement additional security protocols or tiered verification. This example uses four MSBs to identify a QoS level. This is simply a suggested format. Any predetermined bits can be used. It's preferable, however, that the same watermark be used within each packet of the stream. The watermark might not contain a QoS indicator, in which case, all bits allocated for the watermark might be used for a unique identifier, such as that associated with a particular watermark packet key. Figure 1 shows a schematic of how the system routes packets.

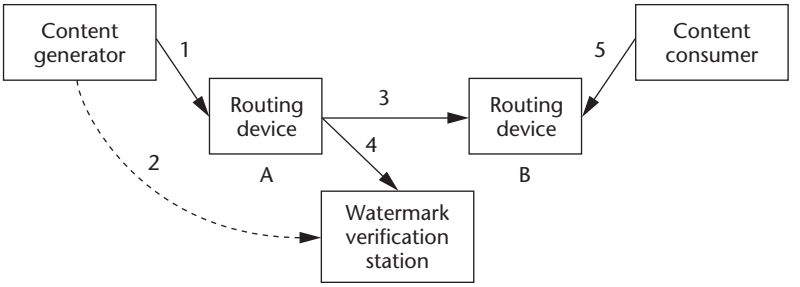

Figure 1. System schematic.

The WID holds all the dependent data. There's only one 32-bit watermark assigned for each stream and one WID created. The watermark packet key may be reused. So the WID might contain a

- 32-bit watermark, inclusive of any QoS indicator,
- watermark packet key,
- hash output from the first block of the flow of data steam,
- hash output from the second block of the flow,
- hash output from the third block of the flow, and
- a series that is bounded by the last block (the flow has a variable length depending on what the data represents).
- hash output from the last block of the flow.

Each router along the flow's path can read the watermark and determine its QoS by using those bits associated with the QoS indicator. Each router can then take appropriate action for prioritizing or deprioritizing each packet. These actions might include choosing a path based on load, reliability, or latency or buffering lower priority packets for later delivery.

The router configuration might enable checking each packet's authenticity. Preferably, the router configuration indicates checking a subset of the packets for authenticity and scaling up to additional cryptographic protocols thereby main-

taining overhead or reducing computational requirements by adjusting security policy consistent with the authentic packet flow. For example, copies of a predetermined, small percentage of watermarked packets might be diverted to a sniffer. Preferably, the sniffer has received the WIDs for all authorized flows either before receiving the flows or in the same time frame. The sniffer compares the watermark of the copied packet to its WID table to find the appropriate WID. If the sniffer doesn't identify a corresponding watermarking key, it deems the packets unauthorized and instructs the router to deprioritize or, preferably, block the flow of the nonauthentic data. If the sniffer finds a corresponding WID, it calculates a hash output for the packet and attempts to match it to the corresponding hash in the WID. If the hash values match, the sniffer instructs the router to permit the flow to continue on its path. If the hash values don't match, the sniffer deems the packets nonauthenticated and notifies the router. Further rules might be associated with any number of scenarios as to why the router has deemed the flow nonauthentic, including notification and reference of the action to a database.

Ideally, the watermark generator software maintains a specific list of sniffers to receive the WID. For each of these, the WID should be sent encrypted and signed, using a public key technology such as PKIX certificates or Open PGP keys. One possible arrangement is having the watermark generator deliver the WID to trading partners who have established a prior business arrangement. The trading partners would pass the WID along to additional devices, eliminating scaling problems on the sender side. These might comprise, moreover, functions handled by the exchange and clearinghouse features.

Generally, it's advantageous for a sniffer to collect twice the original number of bytes to guarantee enough data to calculate a hash, given that the sniffer doesn't know *a priori* the original number of bytes. For large flows, 100:1 ratios might create unacceptably large WIDs. However, as the ratio decreases, the WID delivery channel gets larger. As the ratio increases, the amount of original content necessary to the sniffer increases, as does the amount of the flow that can pass before completion of an authorization check. Making the ratio sensitive to data type and size, or some predetermined policy parameters, dynamically optimizes the system to meet the needs of a particular market. Given this flexibility, overhead will more than likely remain small, compared to more granular accounting and its associated cost savings. Essentially, decisions concerning how much security should be mapped to the flow (for instance, applying a digital signature instead of a hash) are likely to mirror the business models of the markets for which packet watermarking is directed. To more fully extend the benefits of this example, later work will consider additional novel features concerning data management, pricing mechanisms, clearinghouse and dispute resolution methods, and systems.

Conclusions

For electronic networks, any number of data files can occupy bandwidth at some discrete instance of time. The purpose of packet watermarking is twofold. First, it lets bandwidth control devices recognize traffic that should move through the public Internet on specific paths, with either higher-than-normal or lower-than-normal priority. Second, the watermarking lets a bandwidth delivery service monitor its traffic to identify specific content sources. This is for purposes of revenue generation, content or data license management, bandwidth as payment or currency, or any other application where a specific data source needs identification. Watermark sampling requires two pieces of information, or the WID: the watermark key and the labeling information that associates the specific content with a hash array. The distribution of this information requires a secure mechanism because it contains cryptographic material (that is, the watermark key).

Security, like insurance, is a process for managing risk. Cryptographically identifying users demanding packets and subsequently provisioning a particular authenticated path (flow) between users is a basis for enabling bandwidth as currency. Heuristics might be applied as the system learns the best paths for packets to effectively determine subsequent use. Taken to another level, the packets can be further analyzed based on the data's nature, if such identification is available. Packet watermarks and data object watermarks establish responsibility for data's objects or functions (for algorithmic data, such as source, object and executable code). Such responsibility and accountability lie at the heart of a commercially acceptable platform for information commerce. **MM**

Acknowledgments

Thanks to Mike Berry, Peter Cassidy, Nevenka Dimitrova, Yair Frankel, and Rodney Thayer for their helpful contributions to this work.

References

1. G. Simmel, *The Philosophy of Money*, D. Frisby ed., originally published 1907, Routledge, 1999.
2. L. Weschler, *Boggs: A Comedy of Values*, The Univ. of Chicago Press, 1999.
3. A. Danto, *Transfiguration of the Commonplace*, Harvard Univ. Press, 1996.
4. S. Moskowitz, *So This is Convergence?*, Serendip, 1999 (in Japanese), http://www.bluespike.com/papers/convergence.pdf.
5. A. J. Menezes et al., *Handbook of Applied Cryptography*, CRC Press, 1997.
6. *Exchange Mechanisms for Digital Information Packages with Bandwidth Securitization, Multichannel Digital Watermarks, and Key Management,* US Patent Application Serial No. 08/674,726, Patent and Trademark Office, 1996.
7. *Optimization Methods for the Insertion, Protection, and Detection of Digital Watermarks in Digitized Data*, US Patent No. 5,889,868, US Patent and Trademark Office, 1999.
8. *Z-Transform Implementation of Digital Watermarks*, US Patent No. 6,078,664, US Patent and Trademark Office, 2000.
9. L. Lessig, *The Future of Ideas*, Random House, 2001.
10. S. Leibowitz, *Policing Pirates in the Networked Age*, Policy Analysis No. 438, Cato Institute, 2002.
11. J. V. DeLong, "Defending Intellectual Property", *Copy Fights: The Future of Intellectual Property in the Information Age*, Cato Institute, 2002.
12. *Method for Combining Transfer Functions with Predetermined Key Creation*, US Patent Application Serial No. 09/046,627, 1998 (notice of allowability issued).
13. *System and Method For Permitting Open Access To Data Objects And For Securing Data Within The Data Objects*, US Patent Application Serial No. 09/731,039, Patent and Trademark Office, 2000.
14. *Steganographic Method and Device*, US Patent No. 5,613,004, Patent and Trademark Office, 1997.
15. *Method for Stega-Cipher Protection of Computer Code*, US Patent No. 5,745,569, Patent and Trademark Office, 1998.
16. J. Bughin, "Black-Scholes Meets Seinfeld", *The McKinsey Quarterly*, no. 2, 2000, pp. 13-16, http://www.mckinseyquarterly.com/article_page.asp?ar=810&L2=17&L3=66.

Readers may contact Scott Moskowitz at Blue Spike, 16711 Collins Avenue, No. 2505, Miami Beach, Fla. 33160, email scott@bluespike.com.