

Off-Line Viral Economic Systems for Digital Media

Shan He[†], Renan G. Cattelan[‡], Kamal Jain[◇], and Darko Kirovski[◇]

[†] ECE Dept., University of Maryland, College Park, MD, USA

[‡] Instituto de Ciências Matemáticas e de Computação – USP,
São Carlos, São Paulo, Brazil

[◇] Microsoft Research, One Microsoft Way, Redmond, WA, USA

Contact: darkok@microsoft.com

TECHNICAL REPORT MSR-TR-2007-167
DECEMBER 2007

MICROSOFT RESEARCH
ONE MICROSOFT WAY REDMOND, WA 98052, USA
<http://research.microsoft.com>

Off-Line Viral Economic Systems for Digital Media

Shan He[†], Renan G. Cattelan[‡], Kamal Jain[◇], and Darko Kirovski[◇]

[†] ECE Dept., University of Maryland, College Park, MD 20742, USA

[‡] Instituto de Ciências Matemáticas e de Computação – USP,
São Carlos, CEP 13560-970, São Paulo, Brazil

[◇] Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA

ABSTRACT

We propose a novel viral platform for building an incentive-based off-line market for digital media. A user who owns a multimedia clip, can resell it to other end-users such that the revenue is controlled by the copyright owner. For each transaction, the seller retains part of the revenue as an incentive for participating in the distributed economy. Transactions can be committed off-line anywhere, anytime, and by anyone who owns a mobile media player equipped with a near-field wireless. In this paper we propose and discuss a cryptographic algorithm based upon atomic receipt exchange as a foundation of the platform. We also showcase the profit-making potential of the new platform via a network simulator and an economic model extrapolated from real-life data.

1. INTRODUCTION

Most economic ecosystems for digital media are based upon the “on-line store” model that markets, recommends, sells, and stores content onto users’ computing devices. A popular example of such a system is Apple’s combination of an on-line store, iTunes, with a media player device, the iPod [1]. Alternatively, peer-to-peer file sharing systems have gained substantial traction since the introduction of Napster [2]. At its peak in 2001, Napster serviced more than 2.5 billion MP3 downloads per month. iTunes recently announced its billionth transaction since 2001 – a sales rate two orders of magnitude lower than the peak Napster transfer rate. Although hard to compare for numerous reasons, e.g., availability of digital media players substantially differed worldwide for 2001 and 2006, one can observe strong discrepancy in performance.

Both distribution technologies leave a lot to be desired. In a typical file-sharing system copyright owners are isolated from the economic flow, not making any revenues from file trafficking. Consequently, the entertainment industry has sought legal action aimed at financially impairing services that enable file sharing [16]. For certain systems such as BitTorrent [8], media is partitioned and sprayed onto several hosts making it difficult for copyright owners to legally point to a specific copyright violator. Alternatively, the protection of digital content distributed via legal channels such as iTunes, is supported by digital rights management (DRM) tools. Typically, a DRM system encrypts the media using a secret key securely stored in the media player [3]. Such systems suffer from the exposure of the secret key – once revealed, it can be used to arbitrarily edit the copyright data [4, 5, 7]. This problem has affected copyright holders to seek

for revenue on-line primarily via client-server architectures, where majority of the marketing, storage, and processing burden is imposed upon the server farm while limiting customers to purchase clips only when they are on-line.

1.1 An Off-line Economy for Digital Content

In this manuscript, we propose a novel platform for marketing and selling digital content, which enables several key features:

- *off-line sales* – an end-user of a copy of particular digital content can sell the content to another end-user without the immediate assistance of the copyright holder or the service provider.
- *immediate purchase* – pending a successful data transfer, the buyer can play the content immediately.
- *incentive-based sales* – one part of the proceeds is assigned to the copyright holder and another credited to the participating sales-force. Parties’ accounts are updated once either seller’s or buyer’s device establishes a connection to the Internet or some other form of global communication.

The platform aims to enable selling digital content by anyone, anywhere, and anytime – posing few restrictions to the network and business models that can be established within the platform. Initially, two or more devices would discover each other and learn each others’ content via a wireless communication channel, e.g., 802.15.3 WPAN [9]. 802.15.3 supports transfer rates in the 11-55Mbps range, sufficient for music clip transfers in several seconds with a low energy bill.

A buyer and a seller could engage in price negotiation followed by the atomic transaction protocol introduced in this paper. This protocol is fully distributed, off-line, and peer-to-peer; however, the act of committing a transaction, i.e., updating users’ accounts, is based on a client-server architecture. The proceeds of each trade are processed upon connecting seller’s or buyer’s portable media player to a global network such as the Internet or a wireless access point.

Integrity of payments is enforced using two features: a portable tamper-resistant media player and an efficient cryptographic protocol. The protocol relies on public key cryptography [17]. In a simple version, each device stores user’s public-private key-pair locally. The private key is used to sign receipts for transactions performed by the device. This key is protected using strong tamper-resistant hardware [27, 28, 29, 30, 31, 33].




Figure 1: The type of data exchange enabled after “breaking” a protected media player. DarkNet players can exchange files with no limitations - however, when they talk to players in the protected world, they can only do so via the proposed protocol.

Each media player conforms to the transaction protocol that protects the economy as long as the user does not “break” its tamper-resistance [12]. As opposed to traditional DRM systems [4], where by “breaking” one player all players are “broken,” here each device owns a distinct private key. Thus, each player must be “broken” individually. We classify the attacks as follows:

- (i) SOFTWARE-ONLY attacks must be prevented by using a hardware-monitored compartmentalized trusted OS [10].
- (ii) In the PARTIAL HARDWARE ATTACK, the adversary destroys the protected private key while “breaking” a tamper-resistant player [27, 28, 29, 30, 31]. Then, the adversary would alter/replace device’s software to create a device that can locally share data in unlimited fashion. Even though the device is “broken” our objective is still met as owners of “broken” players cannot participate in the viral economy. This type of an attack is hard to prevent fully – thus, we assume that its success should incur relatively high cost for the adversary.
- (iii) In the FULL HARDWARE ATTACK, the adversary gains access to the private key and obtains full access to device’s features. The success of the platform depends upon building a device that can render such attacks cost-ineffective to adversaries. Tamper-resistant devices based, for example, on zeroization mechanisms, have already been built and widely used successfully in similar applications, e.g., electronic wallets [27, 28, 29, 30, 31].

1.2 Related Work

Considering the size of the music market alone estimated at around \$12B in the US, there has been surprisingly few solutions that uniquely address the distribution and economics of digital media. A technology particularly related to our work, and to the best of our knowledge, the first to address incentive-based digital media economies, has been deployed at Weedshare [13, 14]. All sales are executed online as all participants are interconnected to the Weedshare servers during transactions – this greatly limits the economic build-up.

Other types of incentive-based systems have been proposed for peer-to-peer systems with an emphasis on the *free-*

rider problem [15], i.e., the existence of users that participate in sharing files only as consumers, not contributors, thus, increasing contributors’ costs. Nearly all incentive-based peer-to-peer mechanisms are focused on limiting free-riders, who themselves are usually a consequence of the availability of free content on peer-to-peer systems. Free content has recently been greatly reduced in file-sharing systems due to legal action from copyright holders [16]. In the proposed system, much alike Weedshare, we aim at the other part of the content distribution spectrum where copyright holders are not isolated economically from the distribution channels. The goal is to, using the convenience of immediate off-line transactions, sway users from peer-to-peer distribution into another model which directly benefits both copyright holders for improved, inexpensive marketing and scalable e-commerce system; and customers for media availability and economic participation in the distribution chain.

2. ATOMIC OFF-LINE TRANSACTION

There exist four entities in an atomic off-line transaction: seller s , buyer b , service provider p , and trusted authority t . The service provider is contracted by the copyright holders to resell and/or organize the sales of their digital content. The service provider is responsible for realizing the payments in the system via credit cards or other form of banking. Much alike in a traditional e-commerce transaction, the trusted authority issues a public-private key-pair to each entity including certificates that authenticate the distributed public keys. This information is used so that users can authenticate each other and prove identities when buying clips or redeeming credits for transactions.

We assume that RSA is used as a public-key cryptosystem [17] by following the IEEE 1363-2000 standard IFSP- and IFVP-RSA version 2 [18]. For a given entity x , we denote its public-private key-pair as $\{p_x, r_x\}$ respectively. In order to vouch for the authenticity of their public key, each entity other than the trusted authority, owns a certificate $c_x = \{p_x, s_x\}$, which contains the signature $s_x = SP_{r_t}(p_x)$, where function $SP_a(b)$ denotes RSA’s signing primitive of message b using private key a . Certificates are verified by proving $p_x = VP_{p_t}(s_x)$, where function $VP_a(b)$ denotes RSA’s verification primitive of signature b using the public key a . Just as in modern certificate verification protocols, we assume that p_t is known to all devices. Finally, each device is assumed to contain a certificate of the service provider, $c_p = \{p_p, s_p = SP_{r_t}(p_p)\}$, upon enrolling in the service.

2.1 Transaction Objectives

Each atomic transaction must fulfill several objectives related to associated threat models. The basic premise is that either buyer’s or seller’s device is likely to be eventually connected to a global network following an off-line transaction. This way, transactions are eventually committed with p so that b is billed and s is credited with the incentive. To commit a transaction, it is sufficient that only one of the participants connects to p . In such a system, the objective is to prevent manipulations that may benefit either of the entities in an unfair manner. We consider the following threat model:

- A Non-repudiation of executed transactions.** b must not be able to repudiate a transaction after which she downloaded the digital content from s .

- B Mutual initiation.** s must not be able to create an arbitrary transaction with a certain b unless he gains total control over b 's media device either physically or via a software virus. The latter case can be prevented by demanding physical action to initiate a transaction such as a "purchase" button that enables data transmission on contact only.
- C Limited damage in case of device loss.** A lost media device could enable the party who finds it to realize only limited financial gain γ defined by the user. Amount γ equals the total purchasing power that a device may have between two synchronization events with p .
- D Device revocation.** When an adversary misappropriates a media device, she can participate in transactions that are identified by p as fraudulent after the act. Here, after the adversary obtains digital content from s off-line, s and p discover that the transaction was fraudulent when s connects to p . To prevent this, lost or misused devices can be identified, catalogued, and this list distributed to all devices upon connection with p . Thus, an updated seller device should be able to verify the financial validity of the buyer before realizing a transaction.
- E Transaction integrity.** Both buyers and sellers must not be able to alter any information about committed transactions.
- F Robustness to communication failure.** Upon communication failure, a buyer or seller must not be able to enjoy the benefits of the transaction, without all details of the transaction being reflected. For example, b could pay for a media clip and lose connection during download. When connecting with p , b should present her transaction receipt to resume download.
- G Media piracy prevention via traditional methods.** The platform should protect copyright holders from piracy via traditional DRM methods such as symmetric encryption and licenses [3]. Such systems are vulnerable as encryption keys can be reverse engineered from players and decrypted content can be captured either digitally or using an analog recorder.
- H Disabling clients who do not commit transactions.** Certain sellers may refuse to take their sales credits to benefit their "buyers" with free content. A user may certainly decide never to connect its media device online or "break" its device and remove its history of non-committed transactions. In both cases, the user must pay an indirect price: by not being able to participate in the distributed economy and by investing effort and funds into "breaking" the tamper-resistant media device. The goal in this case is to offset the higher likelihood for multimedia piracy due to the convenient wireless transfer between sharing devices, with additional direct and indirect costs incurred by the adversary.

The list of requirements, A-H, resembles off-line usage of credit cards with sellers being able to verify canceled buyers' credit cards upon connection with the issuing bank. The

convenience of the system lies in the effect that a credit card-like payment system is established to support an incentive based economy that benefits all parties involved in the distribution of digital content.

2.2 Transaction Protocol

The set of requirements imposed upon this type of transaction seems equivalent to the constraints involved with simultaneous contract signing [21]. There, two parties who do not trust each other desire to simultaneously commit to a contract over a communication channel. Since simultaneous data exchange is impossible in practice, there exists a need for a protocol which enables this feature. To date, the contract signing problem in the absence of a trusted party has not been solved deterministically and efficiently [22].

We show that a simple protocol for atomic receipt exchange can address the system requirements without the complexity associated with fairness of simultaneous contract signing [21]. In essence, b and s authenticate each other, b sends a signed incentive to buy, s sends a receipt, and only after the acknowledgment of b that she received the receipt, the atomic transaction is executed. Then, s may send the content to b .

I – Authentication and Key Exchange. Initially, the two parties must authenticate each other. This is a task already provided in traditional cryptographic protocols such as TLS1.0 [19]. According to the TLS1.0 Handshake Protocol, the opposing sides perform several tasks:

- exchange certificates, c_b and c_s ; then, each side verifies the opposing side's certificate by proving that $p_s = VP_{p_t}(s_s)$ and $p_b = VP_{p_t}(s_b)$,
- exchange information to compute a 48-byte master-secret used to create session keys,
- establish a way of encrypting and compressing data during the following private communication, and
- establish a session identifier as well as a flag specifying whether the session is resumable.

II – Checking the Revocation List. In order to satisfy requirement D, s must verify whether b has a valid account with p . For that reason, p must synchronize connected players with the latest list of "invalid" players (i.e., their public keys). In order to prevent the list from growing excessively, each account has an expiration date specified in the account's certificate. Players with expired accounts cannot purchase content. Thus, if buyer's account is expired or on the list of revoked players, the transaction is aborted. Otherwise, it proceeds with the buyer's commitment.

III – Marketing. It is important that b receives the content that is marketed. As a marketing ploy, the s may forward to b a version of the content that may be of superior quality compared to the copy that is later uploaded to the buyer. When committing to a purchase, b wants to receive assurance that the clip of interest, a , is of particular identity and quality. There may be several versions of this assurance subprotocol. Here, we outline two examples.

- **III.a – Buyer likes clip, does not know author, title.** In this case, s provides clip's cut-out, a_c , which has been approved by the copyright holder as an advertisement, to b . In addition, the holder provides the purchasing data:

$$m'_1 = \{ID(a), s'_1 = SP_{r_p}(H(a_c, ID(a)))\}, \quad (1)$$

where $ID(a)$ returns a distinct identifier and descriptor of clip a . The descriptor may include clip's coding quality, version, copyright holder, license agreement, and price. Function $H(a)$ returns a cryptographic hash, such as SHA-256 [20], of the clip a . By listening to a_c , computing $H(a_c, ID(a))$, and verifying against s'_1 using provider's public key p_p , \mathbf{b} can get assurance that she will ultimately receive the clip a that \mathbf{p} associated with a_c .

Most importantly, note that \mathbf{s} can keep competitive advantage on the market by not revealing the author and title of the advertised clip to \mathbf{b} . Our system enables this feature – the party who owns a can ask \mathbf{p} to provide $\hat{m}'_1 = \{\hat{ID}(a), s'_1 = SP_{r_p}(H(a_c, \hat{ID}(a)))\}$, where $\hat{ID}(a)$ does not contain identifying information for a . The advertisement receipt \hat{m}'_1 can also be provided to a buyer by a seller. An additional economic tool is system's ability to attach a price to \hat{m}'_1 which a buyer (i.e., potential seller) must pay to obtain. Finally, after purchasing the clip, \mathbf{b} obtains the full $ID(a)$.

- **III.b – Buyer knows author, title, buys clip from seller w/o preview.** In this case, \mathbf{s} immediately sends out:

$$m''_1 = \{ID(a), s''_1 = SP_{r_p}(H(ID(a)))\} \quad (2)$$

to \mathbf{b} who can verify that \mathbf{s} is offering the desired clip without media preview.

IV – Buyer's Commitment. In case \mathbf{b} desires to purchase certain digital content, she commits to the purchase by sending a signed intent of purchase to \mathbf{s} . The intent is represented using $m_2 = \{i, s_2 = SP_{r_b}(H(i))\}$, where $i = \{m'_1 \text{ or } m''_1, \mathbf{b}, \mathbf{s}, P_c\}$ and P_c contains purchase data such as date/time/location,¹ license and price. Message P_c can also include a request to buy an advertisement receipt \hat{m}'_1 for a (see step III.a). The buyer sends m_2 to \mathbf{s} as a transaction request. The seller can verify the purchase intent using buyer's public key p_b . In order for both sellers and buyers to protect their privacy, their public keys p_b and p_s are used as pointers to transaction participants in message i instead of \mathbf{b} and \mathbf{s} for privacy issues).

The price and license may be negotiated between \mathbf{b} and \mathbf{s} . Copyright owners must be careful in setting up pricing rules for their content as buyers and sellers can seek alternative payment channels (e.g., cash, trade). Here is an extreme example. A copyright holder did not assign a minimum price to its content a . The holder relied upon seller's incentives in the form of percentage of revenue to motivate selling the content at a higher price. Sellers could sell a at high price but in cash, circumventing system's payment system. Then, they would report a transaction price of \$0 to \mathbf{p} and retain the full actual revenue to themselves. In order to account for this problem, the copyright holder has to use lower-bounded pricing. The holder has to incorporate this

¹In case transaction participants want to protect their privacy, they should be able to chose whether to record such data within the transaction receipt.

type of “incentive” in its economic model when setting up the price/incentive rules.

Finally, in order to prevent a software attack on the buyer as described in requirement B, the system allows the buyer's commitment to be sent to the seller only upon a hardware-assisted approval by the buyer, e.g., by pressing a “purchase” button on the device.

V – Seller's Receipt. In order for \mathbf{b} to claim her purchase to \mathbf{p} , she must receive a receipt from \mathbf{s} . The receipt is constructed as: $m_3 = \{P_r, SP_{r_s}(H(j))\}$, where $j = \{m_2, P_r\}$ and P_r contains receipt information required by \mathbf{p} . The buyer can verify the receipt using seller's public key p_s . If the verification is successful, \mathbf{b} can claim a from \mathbf{s} or if communication is terminated, from \mathbf{p} . If the latter event occurs, \mathbf{p} can credit the incentive to seller's account even without synchronization with seller's device.

VI – Buyer's Ack. Upon receiving and verifying seller's receipt, \mathbf{b} sends an acknowledgment signal, $m_4 = SP_{r_b}(m_3)$, back to \mathbf{s} . Upon receiving and verifying the acknowledgment, \mathbf{s} can claim the incentive independent of buyer's communication with \mathbf{p} . Hence, \mathbf{b} can commit the transaction with \mathbf{p} independent of \mathbf{s} after step V. For \mathbf{s} to claim his incentives independent of \mathbf{b} , step VI must be finalized.

VII – Content Download. Upon receiving and verifying m_4 , \mathbf{s} starts with the upload of a . The content is encrypted with a session key derived from the session master key (created in step I). The buyer can immediately start enjoying her purchase. If the transaction included the corresponding advertisement receipt \hat{m}'_1 (see step III.a), then \mathbf{s} must upload this data.

The act of downloading the media clip in this protocol is a matter of mutual agreement between \mathbf{s} and \mathbf{b} . The act can be interrupted by lack of power, communication, or intentionally at either one of the devices. The overall transaction is not affected by unsuccessful step VII, as both \mathbf{b} and \mathbf{s} have their receipts to claim the content and incentive independently. The protocol may introduce certain marginal fairness issues discussed in Subsection 2.3. Subsection 2.5 discusses how the protocol can be altered so that the act of media transfer can be probabilistically guaranteed and priced.

VIII – Claiming Incentives. The seller is credited with his incentives upon the following two events.

- The seller received a valid m_4 , in which case he submits to \mathbf{p} the following message $\{m_3, m_4\}$. Upon successful verification of signatures in m_3 and m_4 , \mathbf{p} credits \mathbf{s} with the incentive and forwards the remainder of the revenue to the copyright holder associated with a .
- In the alternate case, \mathbf{b} never received the digital content. When she contacts \mathbf{p} to download the content from its server with a proof of purchase m_3 , \mathbf{s} is credited with his incentive. Both actions are executed pending a successful verification of signatures in m_3 .

Communication failure can occur in steps VI or VII and still the transaction can be committed in the first case by \mathbf{b} and in the second by \mathbf{s} contacting \mathbf{p} . If communication failure or some other form of not conforming to the protocol occurs in steps I–V, the transaction is voided.

2.3 Discover Only, Buy Elsewhere

Interestingly, the communication between \mathbf{s} and \mathbf{b} can be terminated for whatever reason after step V is and step VI

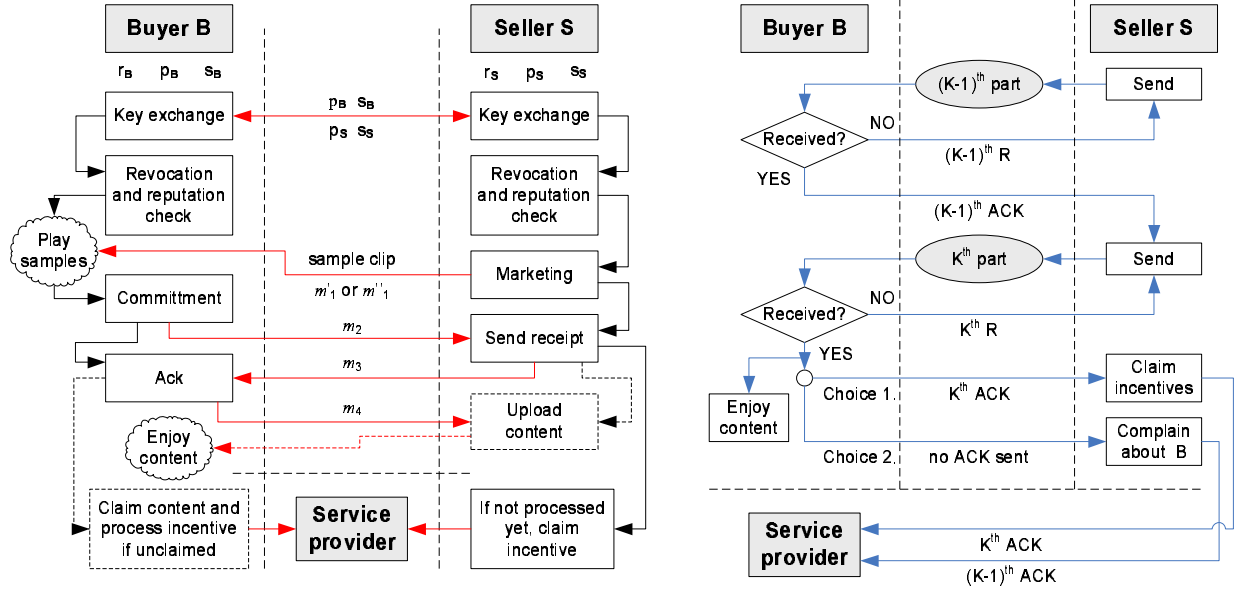


Figure 2: (left) Illustration of steps in the protocol for atomic off-line transaction of digital goods. (right) Events involved in guaranteeing the upload of purchased media to buyer’s media player.

is not finalized. Then, reporting the transaction is up to **b** – if reported, she will obtain the content and **s** the incentive. If she decides not to obtain the content, **s** never receives the incentive as the transaction is never fully completed by **b**. Later, **b** can obtain the same content from some other source. This may be perceived as unfair to the first seller. We refer to this problem as “discover only, buy elsewhere” (DOBE). The DOBE fairness issue can be addressed in our system from two perspectives.

Prevention. First, the protocol in step III.a already prevents DOBE by enabling **s** to advertise a clip without revealing identifying information about it. Since the advertisement receipt \hat{m}'_1 can be priced, in a free market the probability that a buyer walks out of a transaction with info about the clip, is likely to set the price tag on \hat{m}'_1 . Hence, if a seller deems that DOBE is likely in his environment, he can purchase \hat{m}'_1 and use it while marketing.

Handling. Even without an advertisement receipt, **s** can still take certain actions if he suspects **b** will commit DOBE. He may report the receipt of m_2 from **b**. As m_2 contains $ID(a)$, with a simple lookup into its transaction database, **p** can verify whether **b** has bought the same clip elsewhere. Several actions are possible at that point; probably the least costly is to affect buyer’s reputation. Just as in common trading markets such as eBay, one’s reputation is evaluated from a set of positive and negative transaction closures. Subsection 2.4 discusses the maintenance of user reputation in the system. Finally, this approach may have certain negative effects on the system. If users deem their reputation as an important leverage, it is beneficial for sellers to push all their advertisement to prospective buyers hoping that they will be the first seller to offer a clip that a buyer will eventually buy. As buyers can shop for only few clips at once, the only way to prevent this, is to disable push marketing.

2.4 User Reputation

The concept of user (buyer or seller) reputation, common

in on-line marketplaces, can be enforced in our system as well. For example, **p** can maintain reputations by issuing reputation certificates to users upon their synchronization. Each certificate may include a reputation quantifier that can be verified by another party in step II of the off-line transaction protocol. Upon claiming a sale transaction, **p** updates the user with his/her most current reputation status. If the user is a seller, this update must be credited before issuing the transaction credit to the seller. Otherwise, sellers can receive credits and never update their favorable reputation mark. Finally, since a user with a “broken” player can reinstate his favorable mark at will, it is important to set relatively frequent expiration dates on such certificates in order to force users into common updates of their reputation status. The efficacy of such systems in realistic market environments has been a subject of analysis and constant improvement [24, 25] – here, it is beyond the scope of this paper.

2.5 Pricing the Bandwidth

Media download while both devices are off-line, has functional value for both the buyer and the seller. The buyer can play the content immediately. The seller consumes additional energy to transfer a relatively large media file. The seller may chose to avoid uploading the media file to preserve energy as an act of unfairness. For example, with a cost of about US\$5 per communication device, Bluetooth transfers data at a rate of at most 721Kbps, low-cost ZigBee at up to 250Kbps, and more expensive 802.11g devices at 54Mbps. As a common media clip is typically in the 2-8MB range, download can take substantial time and produce a significant energy bill. To address this issue, we propose an additional, optional sequence of steps to the protocol which enables the seller to price the actual download into the transaction. Thus, the buyer can obtain a purchase receipt for one price and both the receipt and the content for another, higher price. The optional part of the protocol is

illustrated in Figure 2(right).

In order to realize such a transaction, a buyer **b** has to specify the type of transaction (receipt or receipt+media) as well as the price when creating the intention to purchase. This is denoted in the field P_c in step **IV**. At step **VII**, a seller **s** partitions the content a into K packets and sends them independently to **b**. One of the objectives is to force **b** to upload all K packets in order to play any perceptually significant portion of a . Thus, **s** initially generates a fresh encryption key k , encrypts a in CBC mode (denoted as $E_k(a)$; [26], pp.229, §7.2.2), and creates a message $e = k || E_k(a)$. Message e is then partitioned into K parts, $\{e_1, \dots, e_K\}$, which are then sent to **b** in decreasing order of their index, i.e., part e_1 is the last, K -th packet sent to **b**. Each packet transmission is followed by an acknowledgment of receipt. The last two acknowledgments, ack_{K-1} and ack_K , in the process are signed by **b**, where $ack_j = SP_{r_b}(H(i||j))$. After receiving ack_{K-1} , **s** sends the last packet e_K . The buyer can decrypt and play the content after this step. However, **b** is still required to send ack_K to **s**. When **s** receives ack_K , he can claim the additional pricing incentive to the service provider by supplying ack_K with all other data as presented in step **VIII**.

Several incident cases may arise in this procedure:

- (i) **b** may receive e_K but fail to send ack_K to **s** due to loss of power or communication. However, **b** can acknowledge the completion of this transaction when she synchronizes with the service provider. Hence, in this case **s** depends upon **b** to communicate eventually with the service provider in order to claim his incentives.
- (ii) After receiving e_K , **b** may maliciously chose not to send ack_K to **s** so that she can obtain the service of downloading the content off-line for free.²
- (iii) **b** may have not sent ack_K because she never received e_K ; **s** cannot distinguish between (ii) and (iii) because communication with **b** has ceased.

The system can address the problem of distinguishing between (ii) and (iii) using at least two strategies. First, users do not decide upon individual protocol actions – in order to be able to alter the protocol steps, **b** must “break” her player’s tamper-resistance and alter its software; two actions that should incur substantial cost. Second, after an incomplete transaction **s** can inform the service provider about the incident. The report includes ack_{K-1} in addition to all other messages described in step **VIII**. Since the likelihood of case (iii) is relatively small, the service provider can affect the reputation of **b** and possibly, additionally charge **b** and credit **s** with his incentive. Thus, user’s reputation becomes a probabilistic reflection of its economic trustworthiness. Even a perfectly policy-obeying buyer is expected to have certain small percentage p of negative feedback. This expectation can be reduced proportionally to the size of e_K , i.e., for that reason, we assume that $e_K = k$. For systems where $p \ll 10^{-2}$, malicious parties can obtain negligible benefits by performing (ii) approximately every $\frac{1}{p}$ transactions. Finally, **s** can report a transaction incident with **b** even though **s** received ack_K – in this case **s** wishes to discredit **b**’s reputation for some reason. To prevent this event,

²Note that **b** still must pay for the purchase receipt in order to download the content.

downloads are always reported by buyers to service providers so that any similar accusations can be cleared.

2.6 Privacy

In any setting where tamper-resistant hardware hosts protected software, typically the issue of privacy is raised. Privacy and security often affect one another and in certain cases it is difficult to ethically resolve and define the rightful balance (e.g., separating crime reporting from privacy protection). In our system, we aim to adopt a common but controversial standard applied in banking and other services where the service provider as a trusted authority keeps record of all transactions in a manner that protects user privacy. With all the ambiguities of such a protection standard, the frontier for privacy protection can be defined from the perspective of the buyer and seller. Ultimately, one would not want that the buyer or seller can show the transaction receipt to a third party in order to reveal seller’s or buyer’s identity respectively, with the associated type of content.

The buyer and the seller exchange identifying information when they establish a secure connection in step **I**. As the public key of either of the users is sufficient to pinpoint its owner, it is important to anonymize user public keys while retaining their full functionality and system security. This can be achieved by distributing single-usage public-private key-pairs to users. A participant in a transaction can optionally use such a key-pair in case she wants to stay anonymous. Such key-pairs are supported with certificates issued by the service provider which can set correct expiration dates and reputation scores.

3. THE MODEL

Economic modeling is certainly not a new art [44]. In economics, a model is a theoretical construct that represents economic processes by a set of variables and a set of logical and quantitative relationships between them. As in other fields, models are simplified frameworks designed to illuminate complex processes. In light of the uncertainties that real-life markets typically experience, most economic models typically leave a lot to be desired [36]. Correspondingly, our objective is not to model absolute but relative values in economic ecosystems built for digital media. We use a simple probabilistic behavioral specification for all economic variables and conclude results via simulation. The goal is to compare the novel viral economy for digital media to existing or hybrid systems while exploring economic tools such as dynamic pricing, various marketing strategies, and end-user behavioral models. In this section we present the key economic variables and the evolution process of the viral structures.

3.1 The Variables

The model has three roles: (i) a *service provider*, **p**, who is contracted by the copyright holder to organize the sales and marketing of digital media a ; (ii) a *seller*, **s**, who is an end-user who has already acquired a copy of a and tries to resell it; and (iii) a *buyer*, **b**, any end-user with potential interest in acquiring a from **p** or **s**. The model allows only one service provider and multiple buyers and sellers. These entities, particularly buyers and sellers, can organize in a highly dynamic fashion, posing almost no restrictions to the network and business models than can be established within the platform.




Figure 3: Illustration of the viral market.

The model is configured using the following parameters:

3.1.1 Server Retail Price $P_{SP_0}(t)$

$P_{SP_0}(t)$ models the retail price that \mathbf{b} pays for each copy sold directly by \mathbf{p} at time t . Time $t = 0$ defines the moment when \mathbf{p} introduces the new clip a to the market. In our system, \mathbf{p} has the power to adjust $P_{SP_0}(t)$ over time to address a potentially varying demand for a . For example, by setting up a high $P_{SP_0}(t)$, \mathbf{p} can sway users away from its servers into the viral network and thus, lower transaction costs. In existing “on-line store” systems such as iTunes, we have $P_{SP_0}(t) = \text{const.}$, although we consider dynamic pricing in such scenarios as well.

3.1.2 Server Viral Price $P_{SP}(t)$

$P_{SP}(t)$ denotes the price that \mathbf{p} asks from \mathbf{s} for each copy sold by \mathbf{s} to another user \mathbf{b} at time t . The service provider enforces $P_{SP}(t)$ by issuing certificates with a desired price. Therefore, it is expected that $P_{SP}(t)$ is a monotonically decreasing function as certificates with higher $P_{SP}(t)$ are not likely to spread in the viral network. Intuitively, this is an expected pricing strategy as the content becomes out-of-date and more common among users.

3.1.3 Buyer’s Reservation Price: $P_B(\mathbf{b}, \bar{t})$

$P_B(\mathbf{b}, \bar{t})$ models the price that \mathbf{b} is willing to pay for a at time \bar{t} . $P_B(\mathbf{b}, \bar{t})$ is generated individually for each user \mathbf{b} . The time axis \bar{t} is shifted w.r.t. t as the pricing curve is initiated from the moment \mathbf{b} learns about a for the first time via traditional or viral marketing. Our $P_B(\mathbf{b}, \bar{t})$ model accounts for the initial appeal that a produces with users, i.e., a user is likely to pay more for a freshly discovered content; this price decays with time.

This model is a direct consequence of the Bass diffusion model that classifies users into innovators, early adopters, early and late majority, and laggards [45]. Data shows that eagerly anticipated media typically has exponentially decreasing diffusion patterns different from the Bass model [46, 47, 48, 49]. Examples of such diffusion patterns are illustrated in Figure 4. Consequently, the question arises on how buyer’s reservation price should be defined to reflect these models [50], in particular for eagerly anticipated media as it is the largest source of profits. To the best of our knowledge, such modeling has not been done to date.

In this paper we propose a simple exponential model for $P_B(\mathbf{b}, \bar{t})$. We estimated its parameters based upon a simple user study on 50 subjects. Its existence in the overall eco-

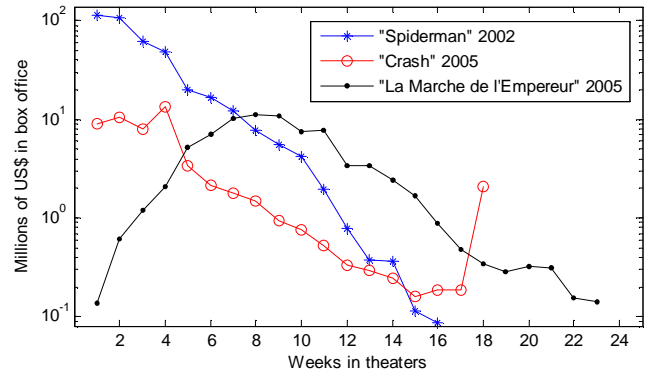


Figure 4: Diffusion patterns at the box-office for three popular movies. There are two examples of eagerly anticipated content as well as one example of a Bass diffusion pattern.




Figure 5: Illustration of the models developed for $P_B(\mathbf{b}, t)$ and $q(\mathbf{s}, I_S)$.

nomic model is only exemplary as its accuracy is only speculative due to the small number of subjects who helped define the model. Figure 5(left) illustrates the $P_B(\mathbf{b}, \bar{t})$ curves generated for $N = 2000$ users. We generated $P_B(\mathbf{b}, 0)$ using a χ^2 -pdf with an expander variable that linearly increases the mean value of $P_B(\mathbf{b}, 0)$. In Figure 5(left), we generated $\frac{1}{N} \sum_{i=1}^N P_B(\mathbf{b}_i, 0) = \1.14 . Finally, we modeled the time-varying decay using a Gaussian pdf – note that the model is such that shortly after discovery, in certain cases $P_B(\mathbf{b}_i, \bar{t})$ may also increase probabilistically for certain buyers. Eventually, buyer’s interest phases out to an asymptotic low price $P_B(\mathbf{b}, \infty)$. We generated $P_B(\mathbf{b}, \infty)$ using the same distribution model as $P_B(\mathbf{b}, 0)$ with a lower mean value; for example in Figure 5(left) we have $\frac{1}{N} \sum_{i=1}^N P_B(\mathbf{b}_i, \infty) = \0.79 .

3.1.4 Seller’s Incentive I_S

I_S denotes the incentive that \mathbf{s} receives for selling a copy of a to \mathbf{b} . We model seller’s behavior independent of t by establishing a probability $q(\mathbf{s}, I_S)$ that a transaction is executed based upon I_S . We generate the $q(\mathbf{s}, I_S)$ model individually for each seller \mathbf{s} – we model the quantifying curve using a Gaussian cdf with a varying mean and low variance that points to relatively sharp, phase-transition-like decision-making by the sellers. Figure 5(right) illustrates the $q(\mathbf{s}, I_S)$ curves generated for 2000 users. In the experiments, we used the $\{P_B, q\}$ -model presented in Figure 5.

3.1.5 Viral Transaction Price P_S

Viral transaction price P_S denotes the price that \mathbf{b} pays to \mathbf{s} for a where $\mathbf{s} \neq \mathbf{p}$. P_S has two components: $P_S =$

$P_{SP}(t) + I_S$. The buyer and the seller engage in price negotiation prior to the transaction. Thus, P_S is established based upon the equilibrium of $P_B(\mathbf{b}, t)$ and $q(\mathbf{s}, I_S)$. The price negotiation algorithm is reviewed later in the text.

3.1.6 Marketing Strategy $R(t)$

$R(t)$ represents the efficacy of traditional marketing for \mathbf{p} . A buyer can learn about a in two ways, by traditional or viral marketing. The key differentiator is the fact that viral does not, whereas traditional marketing does cost the service provider. $R(t)$ is the ratio of users informed via traditional marketing at time t over the entire buyer population. Summation of $R(t)$ over time comprises the overall traditional marketing effort. Typically we can have $\sum_{t=0}^{\infty} R(t) > 1$ due to marketing imperfections, i.e., the effort of marketing may reach already informed users at each time step.

3.1.7 Probability of Transaction Occurrence q_T

If the price is affordable, \mathbf{b} buys a with probability q_T . We use this parameter to model the convenience of buying from anyone, anywhere, and anytime using the proposed off-line peer-to-peer transaction protocol. Client-server scenarios do not have this feature and should be penalized by reducing q_T . Thus, q_T takes two values, q_{T0} and q_{T1} when content is sold by \mathbf{p} or virally respectively ($q_{T0} \leq q_{T1}$).

3.1.8 Non-Recurring Start-Up Cost C_{SU}

C_{SU} models the non-recurring start-up cost that \mathbf{p} pays for initiating the business. It relates to building the server farm, the related software, and other intangibles. In our model C_{SU} is linearly proportional to the maximum number of simultaneous³ server calls that \mathbf{p} receives from its account holders over the content lifetime.

3.1.9 Operational Costs C_{OP_0} and C_{OP_1}

C_{OP_0} and C_{OP_1} model the total operational cost that \mathbf{p} pays for a single direct and viral transaction respectively. The costs include the network bandwidth, maintenance of the server farm, etc. Because reporting a viral transaction does not involve content download, we assume that $C_{OP_0} \gg C_{OP_1}$.

3.1.10 Cost of Traditional Marketing C_M

C_M models the cost of traditional marketing. At a moment t , \mathbf{p} spends $C_M R(t)N$ in order to inform $R(t)N$ randomly selected users from the user population about a . In addition to uninformed buyers, this effort may redundantly inform existing sellers as well as buyers who know about the content but have not purchased it for some reason.

3.2 The Optimization Objective

Regardless of the transaction protocol that builds the economic ecosystem, the objective that \mathbf{p} has, is simple to define. We can compute the profit, Π , fetched by \mathbf{p} as follows:

$$\begin{aligned} \Pi = & \sum_{t=0}^{\infty} [(P_{SP_0}(t) - C_{OP_0})N_{B0}(t)] \\ & + \sum_{t=0}^{\infty} [(P_{SP}(t) - C_{OP_1})N_{B1}(t)] \\ & - NC_M \sum_{t=0}^{\infty} R(t) \\ & - C_{SU} \max_{t=0}^{\infty} [N_{B0}(t) + \alpha N_{B1}(t)]. \end{aligned} \quad (3)$$

where $N_{B0}(t)$ denotes the number of direct buyers from \mathbf{p} at time t and $N_{B1}(t)$ denotes the number of viral buyers at time t . Parameter α is used to scale the setup costs for the case when transactions are executed directly w.r.t. \mathbf{p} (including the downloads) and virally. We simplify the parameter space by assuming that $\alpha = \frac{C_{OP_1}}{C_{OP_0}}$.

Given an existing $\{P_B, q, q_T\}$ user-behavior model, the optimization goal can be defined as:

$$\arg \max_{\substack{P_{SP_0}(t) \\ P_{SP}(t) \\ R(t)}} \Pi. \quad (4)$$

3.3 The Viral Network Model

We model the viral marketing network as a time-varying graph $\mathbb{G}\{\mathbb{N}, \mathbb{E}(t)\}$, where $\mathbb{N} = \{\mathbf{n}_1, \dots, \mathbf{n}_N\}$ is the set of all N users in the network, $\mathbb{E}(t)$ is the set of all bidirectional edges in the network at time t . An edge $e(\mathbf{n}_i, \mathbf{n}_j, t) \in \mathbb{E}(t)$ signifies that users \mathbf{n}_i and \mathbf{n}_j can exchange data at time t .

We developed a model for $\mathbb{E}(t)$ inspired by the work on epidemic data dissemination in complex networks [51, 52, 53]. There, information spreads like an epidemic through local interaction between nodes. After obtaining the information, the nodes forward the message to their neighbors in a random manner. Eventually, the entire system becomes “infected” with information. We rely on the Scale Free (SF) network model which has been shown to model more accurately large scale networks such as the Internet [37]. We denote as k_i the degree of a node \mathbf{n}_i . The SF model distributes node degrees across the network according to the Power-Law: $\Pr[k_i = \kappa] \sim \kappa^{-\gamma}$ with $\gamma \leq 3$ for most real-world networks.

Although the SF network model has been assumed as an accurate model of Internet connectivity based upon experimental data [37, 54], modeling of various types of social networks via the SF network model has been disputed. For example, it has been shown that if we look at communities of interests in a specific topic, discarding the major hubs of the Web, the distribution of links is no longer a power law but resembles more a normal distribution, as observed in [55]. Still, we speculate that the considered viral network should take the shape of a SF network as many collaborative networks have already demonstrated SF properties.

We used the Barabasi-Albert (BA) SF network construction algorithm [37] which mimics the growth of the Internet. It starts with a set of m_0 nodes and a newly introduced node is attached to m existing nodes at each time step. The probability that a link connects the new node to an existing node is linearly proportional to the actual degree of the

³At time step t .

existing node. This procedure is repeated until all nodes are added – the resulting network is expected to follow the Power-Law degree distribution with $\gamma = 3$ and average degree $\bar{k} = N^{-1} \sum_{i=1}^N k_i = 2m$. This network is an example of a highly heterogeneous network where the degree distribution has unbounded fluctuations. An accurate model for building time-varying edges $f : \mathbb{E}(t) \rightarrow \mathbb{E}(t+1)$ is an important component of the overall ecosystem model, however for brevity and lack of experimental data and related work [38, 39] we chose to simplify our presentation. In the current model, we generate $\mathbb{E}(t+1)$ independently from $\mathbb{E}(t)$ while preserving individual node degrees in \mathbb{G} .

3.4 Model Evolution

The model evolves in three steps: (A) traditional marketing, (B) direct sales from \mathbf{p} , and (C) viral sales.

Step A – Marketing. Initially, \mathbf{p} markets the upcoming media clip a by spending C_M to inform a single user. At the end of this stage, $NR(t)$ users are potential buyers of a . The only location they can buy a at this moment is \mathbf{p} .

Step B – Direct Sales: At this stage, \mathbf{p} starts selling a at price $P_{SP_0}(t)$. The informed users initiate their reservation price curves to obtain $P_B(t)$. The potential buyers whose $P_B(t)$ is higher or equal to $P_{SP_0}(t)$ buy a from \mathbf{p} with probability q_{T0} . Each buyer who has purchased a automatically becomes a seller. The remaining buyers who couldn't afford the content wait until the next step.

Step C – Viral Sales: At this stage, each seller in \mathbb{N} contacts its neighbors to virally market a . Contacted buyers who learn about a for the first time initiate their reservation price curves. All informed buyers enter the price negotiation process with their neighboring sellers with probability q_{T1} .

In the simple case when a buyer \mathbf{b} is contacted by a single seller \mathbf{s} , we assume that \mathbf{s} will negotiate a sales price equal to $P_B(\mathbf{b}, t)$. This can be done simply by starting from a relatively high price and slowly reducing the offer. Then, \mathbf{s} calculates his incentive as $I_S = P_B(\mathbf{b}, t) - P_{SP}(t)$ and decides to participate in the transaction with probability $q(\mathbf{s}, I_S)$.

If \mathbf{b} faces a group of K sellers $\mathbb{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_K\}$ at the same time, then \mathbf{b} will probabilistically negotiate the lowest price possible from \mathbb{S} . The buyer is going to offer a low price (e.g., $P_{SP}(t)$), then gradually increase it until one of the sellers accepts it. For each offered price, each seller $\mathbf{s}_i \in \mathbb{S}$ would compute the corresponding incentive I_S and enter the transaction with probability $q(\mathbf{s}_i, I_S)$.

Finally, all informed buyers who were not contacted by \mathbf{p} in step A ponder a transaction with \mathbf{p} . If their current reservation price is higher than $P_{SP_0}(t)$, they make the purchase at that price with probability q_{T0} .

Stages A-C are repeated in the model for each time step t . The model can be stopped until all users buy the clip or when t reaches a particular deadline T .

The presented model and its evolution target realistic scenarios. We pertain to model the artistic appeal that media causes when first experienced. As our platform supports buying media in coffee shops, bars, restaurants, on the beach, at stadiums, essentially at arbitrary locations and times, one can assume that the artistic appeal of media should push users to pay more and/or with higher likelihood shortly after they have been introduced to the media for the first time. While viral marketing has been successful in many economic scenarios, media may be one of its best

targets. For example, historically, for millennia music has been conveyed virally for the lack of any other mechanism. Similarly, viral marketing tends to create sense of community among groups of users who engage in file trading. Both the social effects as well as the unparalleled convenience in managing data are some of the reasons why we speculate that the proposed system is a significantly more powerful economic platform for end-users than the existing “on-line store” model. The benefits that stem from the proposed platform are not unilateral; copyright holders can also adjust their marketing and pricing strategies to optimize profits. In the remainder of this paper, we focus on analyzing copyright holder's strategies for profit optimization and comparing these results with the state-of-the-art.

3.5 The “On-line Store” Model

The “on-line store” model, typical of modern media distribution, can be presented as a subspace in the proposed generic model. By setting the following parameters:

$$P_{SP_0}(t) = \text{const.} \quad P_{SP} = \infty \quad q_{T1} = 0, \quad (5)$$

we do not build the viral network, hence, we allow for sales only from \mathbf{p} . For a given $\{P_B, q, q_T\}$ user-behavior model, \mathbf{p} is aiming to maximize its profit:

$$\begin{aligned} \Pi &= \sum_{t=0}^{\infty} [(P_{SP_0}(t) - C_{OP_0})N_{B0}(t)] \\ &\quad - NC_M \sum_{t=0}^{\infty} R(t) - C_{SU} \max_{t=0}^{\infty} N_{B0}(t). \end{aligned} \quad (6)$$

The evolution of the “on-line store” model is equivalent to the generic case presented in Subsection 3.4. As the viral marketing phase does not exist, stage C is ignored.

The “on-line store” model is important for an additional reason. It can be used to compute system parameters such as C_{SU} and C_M based on well-documented data. For example, an existing survey points to the cost-profit distribution for iTunes [40]. The $P_{SP_0} = 99\text{¢}$ price-tag is split as follows:

- copyright holders receive 60-70¢ as revenues before their operational costs; we assume that the marketing costs of copyright owners may range within 5-10¢,

and on the service provider's side only:

- marketing costs 5-10¢,
- actual financial transaction costs 10-15¢,
- staff payroll ranges from 3-5¢,
- negotiated bandwidth may cost 2-5¢, and
- start-up costs to establish an “on-line store” are estimated at 5-10¢⁴,

which accounts for profits at the service provider's side that may range from a 16¢ loss to a 14¢ gain.⁵ Using the costs

⁴The original estimate in [40] ranges between 2-3¢. Infrastructure costs were amortized over 10 years – however, studies show that server farms are typically refreshed every 3-4 years. Thus, we have increased these costs accordingly.

⁵The reason why the estimated range is relatively wide lies in the fact that Apple has never published an official study for their pricing/cost model for iTunes.

presented here, for an arbitrary number N of users in the ecosystem, we can compute C_M and C_{SU} based upon the cost components represented in Eqn.6. These costs can then be used in simulations of other economic ecosystems.

3.6 The DarkNet Model

In the proposed ecosystem, the DarkNet [11] is reality just as in the “on-line store” model. “Broken” players could engage in unlimited file sharing with other “broken” players, which is DarkNet’s state-of-the-art. Fortunately, “broken” players cannot invent new and replay or adjust existing transactions – thus, players from the DarkNet cannot claim any additional financial benefits from the system. Interestingly, pirated media could be sold via the proposed ecosystem so that the profits belong to the rightful copyright owner (and no other destination) [64].

We model two costs related to the DarkNet that balance the decision for each user whether to make the investment to “break” her player. Let’s assume that the number of copies that the user obtains freely after switching to the DarkNet is K . We denote the cost of “breaking” a device as X . This cost should be substantial (i.e., users should not be able to “break” players via a software attack only). As opposed to early tamper-resistant systems [33], modern e-cash systems based upon secure smart-cards have fared well in adversarial environments [34]. Next, we denote the likelihood that a DarkNet user is identified as ϵ_1 . The legal fine in this case is J . A DarkNet user can now amortize costs to compute the overall price per clip: $P_{D0} = K^{-1}(X + \epsilon_1 J)$. A user \mathbf{u} is likely to switch to the DarkNet if P_{D0} is smaller than the average price \mathbf{u} pays per clip minus the expected average incentives \mathbf{u} receives for selling purchased content.

Here, an important question arises: if a DarkNet user \mathbf{u} can resell⁶ pirated material, can the profits that \mathbf{u} makes help reduce P_{D0} ? We mark this question for further research – obviously, this problem is non-existent if sellers are allowed to sell only legally purchased content.

The price of pirated material in the state-of-the-art file sharing systems equals $P_{D1} = \epsilon_2 J K^{-1}$, where ϵ_2 is the probability that a DarkNet user is identified. One may argue that $\epsilon_1 \ll \epsilon_2$ as it is easier to associate an IP address with user’s correct identity than detect off-line DarkNet users. However, as technology improves and media players with wireless broadband become reality, we expect that eventually $\epsilon_2 = \epsilon_1$. Then, DarkNet’s state-of-the-art may improve even further, leaving few choices to copyright holders beyond the proposal introduced in this paper.

4. DYNAMIC PRICING

From Eqn.4 we observe that the optimization goal for \mathbf{p} is to find out a pricing strategy for both $P_{SP_0}(t)$ and $P_{SP}(t)$ and a marketing strategy $R(t)$ that maximize its profit Π for an unknown $\{P_B, q, q_T\}$ user-behavior model. While most current pricing systems adopt $P_{SP_0}(t) = \text{const.}$ [56], here we propose a novel dynamic pricing heuristic which addresses variable demand. Although studied dynamic pricing methods have shown efficiency in profit optimization [57, 58, 59, 60], we could not deploy previous work due to the novelities introduced in our ecosystem. We developed a novel derivative-following dynamic pricing strategy [57, 61, 62]

⁶Seller \mathbf{s} can sell a pirated clip a in the legal ecosystem only if \mathbf{p} issues a certificate to \mathbf{s} for a .

that addresses the key trade-offs in our economic ecosystem under a fixed marketing plot. The proposed heuristic consists of two parts: one for $P_{SP_0}(t)$; the other for $P_{SP}(t)$. At each time step t , \mathbf{p} adjusts $P_{SP_0}(t)$ and $P_{SP}(t)$ based on the sales from the previous time step $t - 1$. The price adjustment $\Delta(t)$ can be additive or multiplicative. We define two efficiency statistics $\Gamma_{SP_0}(t)$ and $\Gamma_{SP}(t)$:

$$\Gamma_{SP_0}(t) = \frac{N_{B0}(t)}{[N - N_S(t-1)]} \frac{1}{R(t)} \quad (7)$$

$$\Gamma_{SP}(t) = \frac{N_{B1}(t)}{N - N_S(t-1)} \frac{N}{N_S(t-1)}, \quad (8)$$

where $N_S(t)$ denotes the number of sellers at time t . The first term in Eqn.7 measures the ratio of potential buyers who have purchased content in the previous step. Thus, $\Gamma_{SP_0}(t)$ represents the ratio between the actual and expected revenues stemming from transactions directly with \mathbf{p} . Similarly, the first term in Eqn.8 is the inverse of the actual sales fetched virally, while the second term measures the revenue potential stemming from the viral network. We alter P_{SP_0} based upon the sale efficiency derivative:

$$\Delta P_{SP_0}(t) = \Delta P_{SP_0}(t-1) \text{sign} \left[\frac{\partial \Gamma_{SP_0}(t)}{\partial t} \right], \quad (9)$$

where $\theta = |\Delta P_{SP_0}(t)|$ is the magnitude of the dynamic price updates. Parameter θ affects the final profits negligibly if the time sampling unit is of fine granularity.

Algorithm 1 Dynamic Pricing for $P_{SP_0}(t)$ and $P_{SP}(t)$

- 1: $t = 0$: $\Gamma_{SP_0}(0) = \Gamma_0$, $\Gamma_{SP}(0) = \Gamma_1$, set initial prices to $P_{SP_0}(0)$ and $P_{SP}(0)$, initialize $\Delta P_{SP_0} = \Delta P_{SP} = \delta > 0$;
 - 2: $t = 1$: $P_{SP_0}(1) = P_{SP_0}(0)$ and $P_{SP}(1) = P_{SP}(0)$.
 - 3: $t > 1$: calculate $\Gamma_{SP_0}(t-1)$ and $\Gamma_{SP}(t-1)$ according to Eqns.7 and 8 respectively.
- $$\Delta P_{SP_0}(t-1) = \Delta P_{SP_0}(t-2) \text{sign} \left[\frac{\partial \Gamma_{SP_0}(t-1)}{\partial t} \right].$$
- $$\Delta P_{SP}(t-1) = \Delta P_{SP}(t-2) \text{sign} \left[\frac{\partial \Gamma_{SP}(t-1)}{\partial t} \right].$$
- $$P_{SP_0}(t) = P_{SP_0}(t-1) + \Delta P_{SP_0}(t-1),$$
- $$P_{SP}(t) = \begin{cases} P_{SP}(t-1), & \Delta P_{SP}(t-1) > 0 \\ P_{SP}(t-1) + \Delta P_{SP}(t-1), & \text{otherwise} \end{cases}$$
-

The intuition behind this heuristic is that when $\Gamma_{SP_0}(t)$ is higher than $\Gamma_{SP_0}(t-1)$, the price change from time $t-1$ is deemed as efficient and we maintain its direction and vice versa. $\Delta P_{SP_0}(1)$ is determined similarly by comparing $\Gamma_{SP_0}(1)$ with a constant $\Gamma_{SP_0}(0) = \Gamma_0$, which represents the expected sale efficiency at $t = 0$.

New prices in the viral network are introduced by issuing appropriate sales certificates. As higher prices are unlikely to spread in the viral network, we put a constraint on the pricing strategy $\partial P_{SP}(t)/\partial t \leq 0$. Therefore, the pricing strategy for $P_{SP}(t)$ is derived equivalent to Eqn.9 with a difference that it maintains (does not increase) the price when better sale efficiency is observed. The dynamic pricing algorithm is initialized similar to the heuristic for P_{SP_0} . Both heuristics are summarized in Algorithm 1.

As long as the time sampling granularity for this strategy is relatively high compared to the smoothness of the revenues curves, the proposed dynamic pricing algorithm should perform near-optimally [63] as it has the ability to

adapt to the market demand. In this version of the paper, we do not evaluate joint pricing-marketing strategies; instead we assume a specific $R(t)$ empirically optimized for each of the evaluated systems: the off-line viral ecosystem as well as the “on-line store” model.

4.1 Accounting for Certificate Updates

When \mathbf{p} issues a certificate that sets a new asking price for media a , it can reach sellers in two ways. First, all content sold at time t is associated with the new cert. Second, devices that already contain a (i.e., current sellers) can connect to \mathbf{p} and refresh their certs. In the first case, the price-setting cost is included in C_{OP_0} ; in the second one, the cert-update action can be modeled separately as operation cost C_{OP_2} (mainly, the bandwidth costs for cert delivery; typically, certs are updated in bundles). For our model, we assumed $C_{OP_2} = \beta C_{OP_1}$, $\beta = 0.1$.

We propagate certs during system evolution as follows. For each time step t , a ratio ω of all existing sellers in \mathbb{G} connects to \mathbf{p} to update their certs. The overall cost of cert updates is computed as $C_{CU} = \sum_{t=1}^{\infty} \omega N_S(t)$ and it is subtracted from Π in Eqn.3. C_{CU} can be reduced if sellers are allowed to update their certs virally.

Accounting for cert updates results in pricing diversity across the viral network. This is certainly a realistic consequence and manifests with a trade-off. Since $P_{SP}(t)$ is expected to decrease over time, \mathbf{p} actually makes more profit from a single sale when an outdated cert is used. In addition, \mathbf{p} is expected to lose sales volume because higher-priced out-of-date certs are more likely to match fewer buyers’ $P_B(t)$ curves. We have experimented with ecosystems where these effects are both ignored and accounted for.

5. SIMULATION RESULTS

We examine three scenarios for the off-line system: fixed and dynamic pricing and dynamic pricing with cert propagation. For comparison purpose, we also examine fixed and dynamic pricing for the “on-line store” model. We first study the performance of both systems with the same marketing strategy. Then, we vary the marketing for each system to maximize their profits and compare the results. Finally, we look into the effect of content popularity on profit, a feature not exhibited by fixed pricing systems.

5.1 Equivalent Marketing Strategies

We first compare the performance of the viral system vs. the “on-line store” model under a specific traditional marketing strategy: after the initial $R(0) = 0.2$; uniform marketing effort is deployed for the remaining T time steps, i.e., $R(t)$ is chosen such that $\sum_{t=0}^T R(t) = 1$. We examine a relatively large and realistic range for $P_{SP_0} \in [0.7, 2.5]$. For each P_{SP_0} value, we look into $P_{SP} \in P_{SP_0} + [-0.2, 0.2]$ and choose a value for P_{SP} that optimizes the profit. C_{OP_0} , C_M and C_{SU} are chosen as described in Subsection 3.5. The total number of users (potential buyers) in the system is $N = 2000$. For each scenario, we examine $C_{OP_1} \in \{\frac{2}{3}, \frac{1}{15}\} C_{OP_0}$, $q_{T0} \in \{0.1, 0.5\}$, and $q_{T1} = q_{T0} + 0.4$. The range for C_{OP_1} reflects the upper and lower bound on the estimated spectrum of operational costs in the proposed economy with respect to the “on-line store” model [40]. Similarly, we addressed marketing content with different appeal to consumers ($q_{T0} \in \{0.1, 0.5\}$) and modeled the appeal of the same content when purchased anytime, anywhere with

$$q_{T1} = q_{T0} + 0.4.$$

The results are shown in Figures 6-9, where “iT” denotes the “on-line store” model, “OL” stands for our proposed off-line system, “F” and “D” represent fixed and dynamic pricing respectively, and “C” means we consider cert propagation. We analyze the results in terms of profit, total sales, and the expected price each buyer pays for the media. All profit results are reported as the combined profit of the copyright holder and the service provider.

Remark 1: Dynamic > Fixed Pricing Dynamic pricing results in higher profits than fixed pricing for both systems. Figure 6 shows that dynamic pricing gives more profit than fixed pricing both the iT (7.7% profit increase) and OL (25.8%) system. Even with the costs of cert updates included, dynamic pricing still offers substantially higher profit than fixed pricing.

Remark 2: OL > iT Profits for the OL systems are significantly higher than for the iT systems under both pricing strategies. The increase can be as high as 169%. Accounting for certificate updates increases operational costs only by a small amount (less than 5% profit decrease). In addition, note that there were few values P_{SP_0} values that resulted in the best iT system being better than the worst OL system. Distribution of profits and operational, setup, and marketing expenses for configurations that fetch optimal profit for each of the five economies under consideration, is presented as pie-charts in Figure 13(left column).

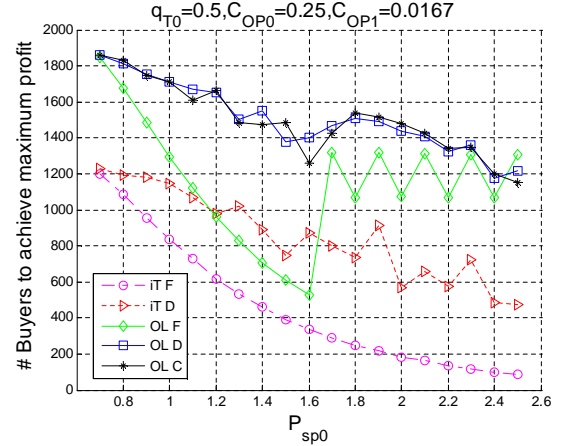


Figure 7: Number of transactions at $q_{T0} = 0.5$.

Remark 3: $q_{T0} \uparrow \Rightarrow \text{profit} \uparrow$ Profit increases as q_{T0} increases. From Figures 6(a,b), one can observe that the partial derivative of the profit vs. q_{T0} is higher for the OL systems. Because OL systems enable users to purchase media near the moment $\bar{t} = 0$ of their introduction to new content, we conclude that OL systems perform better within this dimension of the market. Finally, for the case when $q_{T0} = q_{T1} = 0.5$ one can observe from Figures 6(a,b) that the increase in maximum profit over the searched space $\{P_{SP_0}, P_{SP_1}\}$, is still substantial (over 77%) for OL vs. iT systems.

Remark 4: Improved Customer Satisfaction for OL Systems As shown in Fig.7, dynamic pricing results in higher number of transactions compared to fixed pricing in




Figure 6: Profit: (a) $q_{T0} = 0.1$, $C_{OP1} = 0.017$, (b) $q_{T0} = 0.5$, $C_{OP1} = 0.017$; (c) $q_{T0} = 0.5$, $C_{OP1} = 0.167$.

both systems. In addition, the total number of transactions is significantly higher for the best OL system than for the best iT system. The premier feature behind this consequence is the efficiency of viral compared to traditional marketing. Although this comparison seems marginal with respect to fetched profit, it is important from the perspective of the end-users. Higher number of transactions translates to improved user satisfaction as more users are able to enjoy the media shortly after they were initially introduced to it.

ers start obtaining the content virally, thus, the expected buyer's price decreases. This fact points to an anomaly of the fixed-price iT model that although users pay high price for media, \mathbf{p} still retains low profit due to fewer transactions and high setup costs. This points to the obvious appeal of the DarkNet in this case as high content prices justify the risk of content piracy. Similarly, the incentives given to sellers may sway them away from piracy as their expected payments are lower compared to the iT systems.

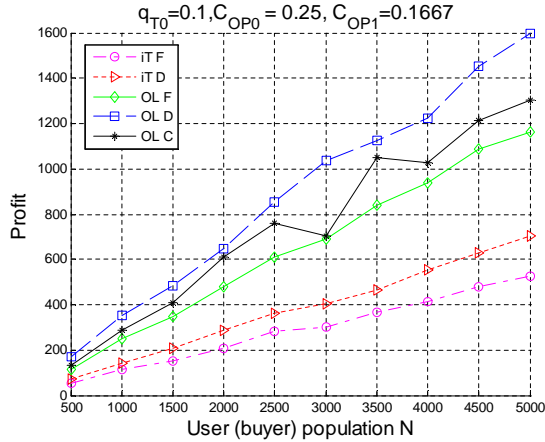


Figure 8: Profit as a function of the user population.

Remark 5: profit $\approx a_1 N + a_2$ We conjecture that profits linearly increase as user population N increases for all variants of both systems. As shown in Figure 8, the rate of increase for the best OL system is more than twice as that of the best iT system.

Remark 6: Expected Buyer's Price With the increase of P_{SP0} , the expected price that a buyer pays does not change in the OL system as much as it changes in the iT systems. Figure 9 illustrates the expected buyer's price with respect to different P_{SP0} values. The expected buyer's price is significantly higher in iT systems with fixed pricing than any other case.

In the fixed-price iT system, buyer's expected price is controlled by \mathbf{p} as buyers have no other choices. Interestingly, buyer's expected price for the fixed price OL system shows a similar trend when P_{SP0} is low. As P_{SP0} increases, buy-




Figure 9: Expected buyer's price with $q_{T0} = 0.1$, $C_{OP1} = 0.0167$.

5.2 Optimized Marketing Strategies

Due to the inherent differences in the iT and OL ecosystems, different marketing strategies result in maximum profits for these models. Here, we fixed the initial marketing expense $R(0) = 0.2$, chose constant marketing for the remainder of the content's sales lifecycle, and varied the marketing intensity to try to maximize empirically the profit for each system. We found that under the same parameter setting as in Section 5.1, $\sum_{t=1}^T R(t) \approx 10$ resulted in maximal profits for the iT model, while $\sum_{t=1}^T R(t) \approx 0$ achieved maximum profit for the OL systems. Figure 11 illustrates the obtained profits under these marketing strategies.

Remark 7: Low Marketing Costs in OL To achieve maximal profit, the iT model needs to put intensive effort on

traditional marketing, while the OL model almost does not need any marketing after the initial $R(0)$. From Figure 11, one can observe that the maximal profit obtained by the OL model is 75% higher than that of the iT model. Distribution of profits and operational, setup, and marketing expenses for configurations that fetch optimal profit for each of the five economies under consideration, is presented as pie-charts in Figure 13(right column).

5.3 Content Popularity vs. Seller Behavior

Here we model media popularity by linearly increasing $\beta P_B(0)$, the average starting buyer reservation price, using a variable β denoted as “expander.” Setting $\beta = 1$ corresponds to the P_B model illustrated in Figure 5. In addition, we model as p the incentive I_S for which the average $q(I_S)$ reaches probability 0.5. The higher the p , the less likely the seller closes on a transaction for a particular incentive I_S . We simulated and recorded the best profits that each of the “expander” and p values produced w.r.t. the remainder of the parameter space – the results are shown in Figures 10 and 12(a-c) and for both iT and OL systems.

Remark 8: $\frac{\partial \Pi}{\partial \beta}_{OL} > \frac{\partial \Pi}{\partial \beta}_{iT}$ and $\frac{\partial \Pi}{\partial p}_{OL} > \frac{\partial \Pi}{\partial p}_{iT}$ The first derivatives of the profit vs. popularity and profit vs. seller functions are higher for the OL systems compared to the iT systems. Figure 12 shows that OL systems consistently give higher profits than iT systems for media of different popularity under various sellers’ behavior patterns. For better comparison, we incorporate the upper and lower bound of Figures 12(a-c) into Figure 10, where we observe that the proposed system has higher profit for content with various popularity than iT system even when the sellers are quite greedy (p as large as 2). The advantage of the proposed system over the iT system is retained for various content and sellers’ behavior patterns.




Figure 10: The effect of content popularity and seller behavior on profits of iT and OL models.

5.4 Why Piracy is Less Likely in OL Models?

According to the piracy model introduced in Subsection 3.6, one could analyze the appeal of the DarkNet. The key statistic to analyze is the buyer’s expected price (Remark 5). From the experiments, we can conclude that the iT system with fixed price is the most prone to piracy as the expected

buyer’s price follows linearly the increase in P_{SP_0} . On the other hand, OL models outsource operating and marketing costs to and share part of the profit with end-users in the form of incentives. Due to the improved availability (i.e., marketing) of content, users can buy content at the peak of their $P_B(t)$ curves generating more revenues per transaction. In our experiments, we actually ignored the effect that users are expected to pay more in the OL model knowing that their expected price will be lower due to the future incentives. In summary, due to reduced operating costs and better availability of content our simulations indicate that the OL systems should fare better with the DarkNet model discussed in this paper. One way for the iT model to address the problem of piracy is to resort to dynamic pricing which has demonstrated significantly better performance than fixed pricing both in terms of profits as well as the expected buyer’s price.

Computing the costs of piracy P_{D0} and P_{D1} in the OL and iT models respectively is not simple as it is difficult to assess all related parameters. Assuming $\varepsilon_2 \rightarrow \varepsilon_1 \approx 0$ as a trend in building portable media players, one can rely only upon the cost X of “breaking” the tamper-resistance of the portable media players to balance the cost of piracy with the benefits of participating in an economy for digital media. From such a perspective, given large enough X , the proposed system is a promising alternative to traditional economic models for distribution of multimedia.

6. DISCUSSION

The economic model and the simulation methodology presented in this paper are important from two perspectives. First, they represent a methodology where hypotheses on business strategies related to the economics of multimedia can be verified and optimized. Second, it opens up the argument on improving the accuracy of similar models. This argument would involve further research on fine-tuning the deployed submodels such as the $P_B(\mathbf{b}, t)$ or the $q(\mathbf{s}, I_S)$ model or the time-variance of the scale free network model. In this paper, we used simplified representations of most submodels extrapolated based upon existing literature [42, 41] and simple user studies. While its accuracy can be disputed, the overall modeling paradigm presented in this paper is the first step towards quantifying the economic efficacy of different multimedia distribution platforms. Finally, the proposed platform could be used for relative as opposed to absolute economic measurements – it could be used to explore a large space of economic situations where competing technologies could be compared in relative manner. With such an objective in this paper – we produced a relative comparison of several economic ecosystems and demonstrated the benefits of viral marketing, anytime/anywhere transactions, and dynamic pricing to copyright holders, service providers, and end-users.

7. SUMMARY

We have introduced a conceptually novel platform for building economic ecosystems for digital goods. Based upon a simple cryptographic protocol, the platform enables users to sell the digital content they own to others so that resulting revenues are controlled by the copyright holders. As a driving force to marketing and sales, sellers retain portion of the revenue as an incentive. To showcase the economic

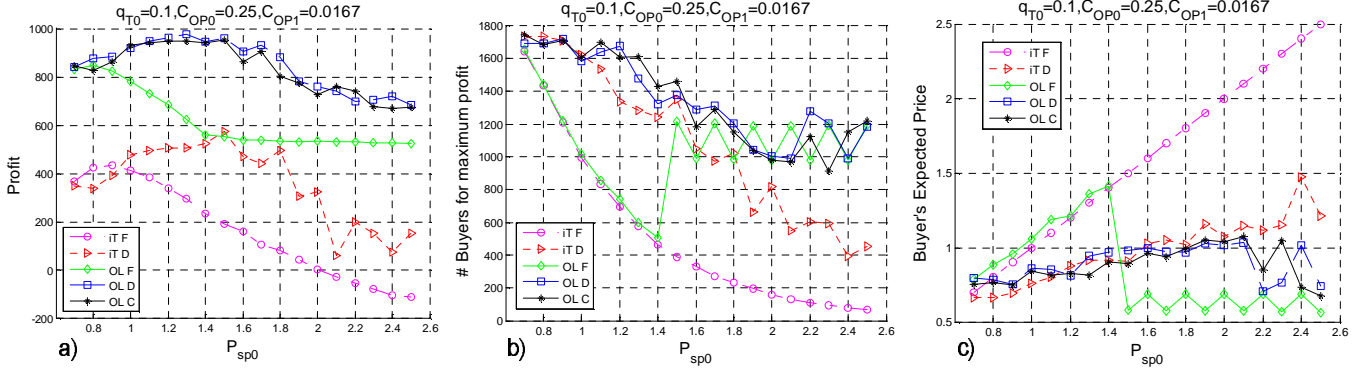


Figure 11: Comparison of the iT and OL models with marketing strategies that aim to maximize their profit. (a) profit; (b) number of total buyers; (c) buyer's expected price.




Figure 12: The effect of content popularity and seller behavior on profits for iT and OL models with (a) fixed pricing, (b) dynamic pricing, and (c) dynamic pricing with cert-update.

benefits of the proposed content distribution paradigm, we introduced a novel economic model based upon a simulator for scale-free networks and viral traffic to compare the novel ecosystem with the existing “on-line store” model. The derived model and simulation methodology are important from two perspectives. First, they represent a methodology where hypotheses on business strategies related to the economics of multimedia can be verified and optimized. Second, they opens up the argument on improving the accuracy of similar models. This argument would involve further research on fine-tuning the deployed submodels such as the $P_B(\mathbf{b}, \bar{t})$ or the $q(\mathbf{s}, I_S)$ model or the time-variance of the scale free network model. In this paper, we used simplified representations of most sub-models extrapolated based upon existing literature [42, 41] and simple user studies. We use the proposed model for relative economic measurements – we explored a large space of economic situations where competing technologies could be compared in relative manner. As a result, the new distribution platform showcased substantial profit increase compared to the “on-line store” model mainly due to substantial reduction in operational and marketing costs, end-user-driven dynamic pricing, and increased sales.

REFERENCES

- [1] Apple iTunes. On-line at: <http://www.apple.com/itunes>.
- [2] S. Fanning and S. Parker. Napster. June, 1999.
- [3] Microsoft Windows DRM. On-line at: <http://www.microsoft.com/windows/windowsmedia/drm>.
- [4] D.S. Touretzky. DeCSS. On-line at: <http://www-2.cs.cmu.edu/~dst/DeCSS>.
- [5] The Hymn Project. On-line at: <http://hymn-project.org>.
- [6] A. Huang. Adding Digital Audio I/O to the Sony D141 CD Player. On-line at: <http://www.xenatera.com/bunnie/proj/cdhack/cdhack.html>.
- [7] S.A. Craver, et al. Reading between the lines: Lessons from the SDMI challenge. USENIX Security Symposium, 2001.
- [8] B. Cohen. Incentives Build Robustness in BitTorrent. Workshop on Economics of Peer-to-Peer Systems, 2003.
- [9] IEEE 802.15 WPAN Task Group 3 (TG3). Available on-line at: <http://www.ieee802.org/15/pub/TG3.html>.
- [10] T. Garfinkel, M. Rosenblum, and D. Boneh. Flexible OS support and applications for trusted computing. HotOS-IX, 2003.
- [11] P. Biddle, P. England, M. Peinado, and B. Willman. The darknet and the future of content distribution. ACM Workshop Digital Rights Management, 2002.
- [12] Federal Information Processing Standards. Security Requirements for Cryptographic Modules. Publication FIPS PUB 140-2, 2002.
- [13] Weedshare Inc. On-line at: <http://www.weedshare.com>.
- [14] Y. Yaacovi. Redistribution of Rights-Managed Content. US Patent (no.) 7,249,107.
- [15] E. Adar and B. Huberman. Free riding on Gnutella. First Monday, Vol.5, no.10, 2000.
- [16] The Digital Millennium Copyright Act of 1998.

- U.S. Copyright Office Summary, 1998. On-line at: <http://www.copyright.gov/legislation/dmca.pdf>
- [17] R.L. Rivest, A. Shamir, and L.A. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, Vol.21, no.2, pp.120–126, 1978.
- [18] IEEE 1363-2000. Standard specifications for public key cryptography. IEEE, 2000.
- [19] T. Dierksa and C. Allen. The TLS Protocol Version 1.0. Internet draft, 1999. On-line at: <http://ietf.org/rfc/rfc2246.txt>.
- [20] Federal Information Processing Standards Publication 180-2. Specifications for the secure hash standard. August 2002.
- [21] M. Ben-or, O. Goldreich, S. Micali, and R.L. Rivest. A fair protocol for signing contracts. *IEEE Transactions of Information Theory*, Vol.36, (no.1), 1990.
- [22] S. Even and Y. Yacobi. Relations among public key signature systems. Technical report, no.175, Computer Science Department, Technion, Israel, 1980.
- [23] J. Goldstein, J.C. Platt, and C.J.C. Burges. Indexing High-Dimensional Rectangles for Fast Multimedia Identification. Microsoft Research Technical Report MSR-TR-2003-38, 2003.
- [24] G. Doukidis, N. Mylonopoulos, and N. Pouloudi. Building Trust Online: The Design of Robust Reputation Reporting Mechanisms in Online Trading Communities. *Information Society or Information Economy? A combined perspective on the digital era*, Idea Book Publishing, 2003.
- [25] C. Dellarocas. Efficiency through feedback-contingent fees and rewards in auction marketplaces with adverse selection and moral hazard. *ACM Conference on Electronic Commerce*, 2003.
- [26] A.J. Menezes, et al. *Handbook of Applied Cryptography*. CRC Press, 1996.
- [27] nCipher Technologies Inc. On-line at: <http://www.ncipher.com/technologies>.
- [28] SafeNet Inc. On-line at: http://www.safenet-inc.com/products/tokens/products_sc_330.asp.
- [29] Infineon Technologies AG. On-line at: <http://www.infineon.com>.
- [30] Axalto Inc. On-line at: <http://www.axalto.com/index.asp>.
- [31] Gemplus Inc. On-line at: <http://www.gemplus.com/smart/rd/publications>.
- [32] G.R. Newman and R.V. Clarke. *Superhighway Robbery: Preventing e-commerce crime*. Willan Publishing, 2003. On-line at: <http://www.eurim.org.uk/activities/ecrime/scale.doc>.
- [33] R. Anderson and M. Kuhn. Tamper resistance – A cautionary note. *Usenix Workshop on Electronic Commerce*, pp.1–11, 1996.
- [34] J. Paynter and P. Law. An arms length evaluation of Octopus. University Of Auckland, Department of Management Science and Information Systems, working paper, 2005.
- [35] BBC News. Digital music revenue triples. October 3, 2005. On-line at: <http://news.bbc.co.uk/1/hi/entertainment/music/4304466.stm>.
- [36] M. Morishima. *The Economic Theory of Modern Society*. Cambridge University Press, 1976.
- [37] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, Vol.286, pp.509, 1999.
- [38] J. Leskovec, J. Kleinberg, C. Faloutsos. Graphs over Time: Densification Laws, Shrinking Diameters and Possible Explanations. *ACM Knowledge Discovery and Data Mining*, 2005.
- [39] J. Leskovec, L. Adamic, and B. Huberman. The Dynamics of Viral Marketing. *ACM Conference on Electronic Commerce*, 2006.
- [40] M. Peitz and P. Waelbroeck. An Economist’s Guide to Digital Music. Work in progress, 2004. Available on-line at: <http://www.gesy.uni-mannheim.de/dipa/32.pdf>.
- [41] T.V. Krishnan, F.M. Bass, D.C. Jain. Optimal Pricing Strategy for New Products. *Management Science*, Vol.45, no.12, pp.1650–1663, 1999.
- [42] B. Robinson and C. Lakhani. Dynamic price models for new product planning. *Management Science*, Vol.21, pp.1113–1122, 1975.
- [43] S. He, R.G. Cattelan, D. Kirovski, and K. Jain. Off-line viral economies for digital content. Microsoft Research Technical Report, 2004.
- [44] H.R. Varian. How to Build an Economic Model in your Spare Time. *American Economist, Passion and Craft: Economists at Work*, University of Michigan Press, 1997.
- [45] F.M. Bass. A New Product Growth Model for Consumer Durables. *Management Science*, Vol.15, pp.215–227, 1969.
- [46] W.W. Moe and P.S. Fader. A Joint Segmentation Model of Consumers and Products Applied to the Sales of Music Albums. The Wharton School, University of Pennsylvania, 1998.
- [47] M.S. Sawhney and J. Eliashberg. A Parsimonious Model for Forecasting Gross Box-Office Revenues of Motion Pictures. *Marketing Science*, Vol.15, no.2, pp.113–131, 1996.
- [48] R. Furukawa, H. Kato, and M. Yamada. A Conceptual Model for Adoption and Diffusion Process of A New Product. *Review of Marketing Science Working Papers*, Vol.1, working paper 3, 2002.
- [49] M. Yamada, R. Furukawa, and M. Ishihara. A Classification Method of Diffusion Patterns with a Class Map. *Acta Humanistica et Scientifica – Kyoto Sangyo University*, Vol.28, no.2, Social Science Series, no.14, pp.59–82, 1997.
- [50] M. Yamada. Inter-class Price Setting in New Product Introductions. *Marketing Science*, 1999.
- [51] Y. Moreno, M. Nekovee, and A. Vespignani. Efficiency and Reliability of Epidemic Data Dissemination in Complex Networks. *Physical Review*, Vol.E69, 055101(R), 2004.
- [52] Y. Moreno, M. Nekovee, and A. Pacheco. Dynamics of Rumor Spreading in Complex Networks. *Physical Review*, 2005.
- [53] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez and D.-U. Hwang. Complex networks: Structure and dynamics. *Physics Reports*, Vol.424, pp.175–308, 2006.
- [54] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. *SIGCOMM*, 1999.
- [55] D.M. Pennock, G.W. Flake, S. Lawrence, E.J. Glover, and C.L. Giles. Winners don’t take all: Characterizing the competition for links on the web. *Proceedings of the National Academy of Sciences*, Vol.99, no.8, pp.5207–5211, 2002.
- [56] M. Bichler, et al. Applications of flexible pricing in business-to-business e-commerce. *IBM Systems Journal*, 2002.
- [57] H.R. Varian. Differential pricing and efficiency. *First Monday*, Vol.1, (no.2), 1996.
- [58] A. Gupta, D.O. Stahl, and A.B. Whinston. The economics of network management. *Communications of the ACM*, Vol.42, (no.9), pp.57–63, 1999.
- [59] P. Dasgupta and R. Das. Dynamic Pricing with Limited Competitor Information in a Multi-Agent Economy. *International Conference on Cooperative Information Systems*, Springer-Verlag, 2000.
- [60] R.M. Weiss and A.K. Mehrotra. Online Dynamic Pricing: Efficiency, Equity and the Future of E-commerce. *Virginia Journal of Law and Technology*, Vol.6, (no.11), 2001.
- [61] G. Gallego and G. van Ryzin. Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizons. *Management Science*, Vol.40, (no.8), pp.999–1020, 1994.
- [62] J.M. Dimicco, A. Greenwald and P. Maes. *Dynamic Pricing*

Strategies under a Finite Time Horizon. ACM Conference on Electronic Commerce, pp.95–104, 2001.

- [63] J.A. Nelder and R. Mead. A Simplex Method for Function Minimization. Journal of Computing, Vol.7, pp.308–313, 1965.
- [64] D. Kirovski and K. Jain. Off-line Economies for Digital Media. ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video, 2006.




Figure 13: Profit and cost distribution for optimum profit points identified in Figure 6 (left column) and Figure 11 (right column).