# Electronic Commerce and the Street Performer Protocol

John Kelsey   Bruce Schneier
{kelsey,schneier}@counterpane.com
Counterpane Systems, 101 East Minnehaha Parkway, Minneapolis, MN 55419

## Abstract

We introduce the Street Performer Protocol, an electronic-commerce mechanism to facilitate the private financing of public works. Using this protocol, people would place donations in escrow, to be released to an author in the event that the promised work be put in the public domain. This protocol has the potential to fund alternative or "marginal" works.

## 1 Introduction

Consider a world without copyright enforcement. People write books, music, etc., but they get paid only for a single performance or print run. Once the work is released, anyone who likes it may make copies and distribute them. In that world, high-quality, easily copied works like stories, novels, reference books, and pieces of music are, in the economic sense, a "public good." That is, the creators of these works must spend scarce resources producing them, but they do not reap most of the benefits.

This leads to the prediction that these works will be produced a good deal less in that world than in ours, and a good deal less than the consumers of these works would like. However, for various technical reasons, we appear to be heading into a world that will look a lot less like our world, and a lot more like that world with no copyright enforcement.

In this paper, we consider a very simple and common approach to funding the production of public goods such as advertisement-free radio and television stations and impromptu music performances in public places. The artist offers to continue producing their freely-available creations so long as they keep getting enough money in donations to make it worth their while to do so. We discuss social, financial, and technical arrangements that can make this approach work fairly well, though we don't believe it will ever provide a complete solution to the problem of paying creators for their creations. We primarily discuss the way a specific instantiation of this idea, called the "Street Performer Protocol," might work.

In the remainder of this paper, we discuss why we believe a continuation of the current situation in copyright enforcement, extended through technical means, is unlikely to work well, how to build the social, financial, and technical arrangements to make this approach work, and the likely attacks on the system. We finish by considering the large number of open questions about this and related schemes.

## 2 Why Copyright Will Be Hard to Enforce in the Future

Before we discuss in detail how our protocol will work, we want to explain why we are so pessimistic about copyright enforcement in the relatively near future. Our pessimism comes from two key beliefs.

First, enforcing copyright laws is made easier when the creation and distribution of high-quality copies of information is relatively expensive and cumbersome. A plant that presses out pirated CDs and a network of trucks and salesmen that distribute them is relatively difficult to hide. Once found, there is no doubt in anyone's mind that the pirates were doing something illegal. Finally, the loss of the expensive equipment and the destruction of the distribution network probably represents a real benefit for the copyright holders, by eliminating a noticeable fraction of the total pirated CD output.

The technology is moving to change all that. Perfect digital copies don't degrade over time, and they take relatively inexpensive equipment to use. A distribution network is already available, in a simple form, today—the Internet. Between the Internet (along with things like e-mail encryption software, anonymous remailers, and the proposed "Eternity Service") and new storage technologies like DVDs, a future pirate is likely to require very little money to get started, and is likely to be an amateur sharing or giving away copies rather than a person making a lot of money running a CD pirating operation.

Our second reason is that the mechanisms for enforcing copyright automatically require a lot of police-state measures. Traitor-tracing schemes require that everyone who buys any copyrighted work provide an ID, and probably ensure that a database of all copyrighted works bought or borrowed from a library by a given person is kept or can quickly be built. Technical enforcement measures can also be used to limit distribution of some writings. Much of the recent activity to prevent copyrighting has amounted to lobbying congress for Draconian anti-piracy laws, laws that limit research into computer security and cryptography, and for laws that seriously restrict what kind of recording and computer equipment is made available for sale to the public.

## 2.1 Technical Solutions: Copyright Commerce Systems

Technical solutions to the problem have been proposed in many places. These tend to fall into two categories:

1. Some schemes attempt to keep the content encrypted except when it is inside a secure perimeter of some kind. The secure device plays, displays, or executes the content only when it is authorized to do so. We will call these "secure perimeter schemes."

2. Some schemes require purchasers of the content to provide some kind of identification, and then embed this identification into the content in some hard-to-remove way. In this case, publication on the net of this content implicates the purchaser, who is probably sued for enormous damages to the copyright holder's intellectual property. (In nearly all cases, the intent will be to deter others from violating intellectual property in the future, rather than to recover losses.) We will call these "traitor tracing" schemes.

Note that it's quite possible to combine both kinds of scheme in the same system.

### 2.1.1 Secure Perimeter Schemes

There are several problems with secure perimeter schemes. The most fundamental problem is simply that, for graphics, video, audio, and text, the value in the content must actually be displayed in a way that the user can see or hear it. (Executable content can be used without leaving the secure perimeter at all, so this argument doesn't apply to it.) This means that, even in an ideal world with impossible-to-break tamper resistance and special sealed devices for all copyrighted materials, making unauthorized copies of this kind of content is still possible. With some custom-designed equipment, it can probably be made fairly easy. (Music is probably easy enough to copy, at some small loss in fidelity, by playing the same piece of music many times, oversampling and re-recording the output, and then processing the result to clean it up as much as possible. Copying video output from a display screen, while clearly possible, looks to be quite a bit more difficult.)

There is also an economic problem. The customer does not see much economic value in purchasing a secure perimeter. Selling a tamper-resistant device useful only to play copyrighted content (e.g., a sealed box with speakers and a video screen) seems difficult. It's easiest to sell or give away software that provides the secure perimeter in which the copyrighted material is kept. However, this is also fairly easy to defeat, even for relatively unsophisticated attackers. (When the attack is finished, it can probably be posted to the Internet.)

Harder, but still reasonable is to sell or give away a special tamper-resistant box that connects to the user's PC or TV, and decrypts copyrighted content when authorized to do so. The satellite and cable TV industries have given us several examples of this, and their record of resisting attack doesn't give us lots of hope for the future of this approach. Along with the various attacks on the box, there are also more general attacks possible–capture the output intended for the screen or speakers, and save it for posting to the Internet. Again, we expect to see software to do this posted to the Internet, as well. (Note that this class of attack hasn't generally been tried on cable and satellite TV systems, because of lack of available bandwidth and storage capacity, and the existence of other, easier attacks.)

Also note that, if the content exists in many forms (e.g., standard music CD, broadcast audio signal, or encrypted audio downloadable from the Internet), the attacker can always save the music from a CD (purchased with cash) to an audio file, and post that file to the Internet. This can be prevented only by never allowing copyrighted music outside these secure perimeters. This involves, among other things, never broadcasting music or films, since they could then be recorded and posted.

With tools like anonymous remailers and the Eternity service, material that's ever posted simply cannot be erased, short of destroying the whole service. This means that one posting of a copyrighted piece of music, video, or text makes it available for free (or at least very cheaply) via the Internet. Indeed, even without the Eternity service, information that is ever posted or made available on the net is very hard to erase, though legal threats can probably get it taken off the major search engines.

### 2.1.2 Traitor Tracing Schemes

Traitor tracing schemes attempt to trace the person who posted the copyrighted material, and to hold him responsible for the losses of the creator of that material. Since these losses are likely to be very large, and since criminal as well as civil penalties may apply, this may deter the person from violating copyright in the first place. This has the additional advantage that it can be implemented entirely within contract law, by requiring anyone who buys copyrighted material to sign a contract agreeing to be liable if his copy of the material is leaked.

The first problem we see with this approach is that it requires the buyer of the copyrighted content to accept the risk that he might be ruined or jailed, if he is accused of posting copyrighted material. He may not have any good reason to trust that the traitor tracing system will get things right. Even if the traitor tracing scheme works, the surrounding system (linking embedded serial numbers or whatever else to human identities) might be subject to attack.

Even worse, a record company or publishing house has relatively little direct incentive to worry about getting the right person. To deter future infringement, they need to make a highly visible example of someone. If it's the right person, so much the better. However, most of the people being deterred by his example will have no idea whether he's guilty or innocent, so the deterrent effect is essentially the same. The record company or publishing house will presumably try to get the right person, but their only financial incentive for doing so is to eliminate one more copyright violator, and to avoid costly lawsuits from the falsely accused person.[1]

In a world in which copyrighted material, once posted, drops a great deal in value, it's probably not possible to hold the copyright violator responsible for most of the loss. He will generally just not have the money. Furthermore, very few personal computers or homes are defended well enough to justify having information inside which, if posted anonymously to the Internet, will cost their owner even a few thousand dollars, let alone millions of dollars. (For comparison, the reader may consider whether he would be willing to keep a briefcase with even $10,000 belonging to his boss in his house, with no additional security or insurance.)

The second problem with this approach is that it requires that every purchaser of copyrighted material present an extremely hard to forge identification. As noted above, this is required for every kind of media, not just for downloading digital content over the Internet; otherwise the smart attacker just buys a CD with cash, loads it onto his computer, and posts it to the net anonymously. These hard to forge IDs must be ubiquitous, and probably end up having to be tied to some kind of national ID card. A determined attacker can try to forge an ID, or can convince some gullible or desperate person to buy the content for him.

The third problem with this approach is that it almost certainly ends up requiring a database somewhere of every piece of copyrighted information anyone has ever purchased. In a world in which nearly all books, movies, and music are purchased online, this creates a really unpleasant destruction of personal privacy. It also raises some interesting questions. Will governments hold this information? How about large media corporations? Will the database records be subject to subpoena by divorce lawyers and independent prosecutors? Will advertisers be able to buy lists of who purchased which book for marketing reasons? What about the security of this database? (How much is the list of everyone who bought *The Satanic Verses* worth on the open market?)

## 2.2 Legal Solutions

Legal enforcement of existing or new copyright laws is made enormously harder by the Internet and other

[1] This incentive problem occurs in many other situations in enforcing laws, e.g., the Olympic Park bombing, and ATM fraud in the UK.

new communications technologies. These technologies allow information to be shared freely, even when governments would rather not have it be shared. This applies as much to copyrighted materials as it does to any other information.

The fundamental enforcement problem with the new technologies is that, in the near future, nearly anyone with a computer and an Internet connection will be capable of posting copyrighted materials to the Internet. These materials, once posted, will be retrievable by nearly anyone, and even without a working Eternity Service, will be quite hard to take off the net once they're put on.

This probably leaves copyright enforcement in the position of spending tens of thousands of dollars of police and court resources shutting down each copyright violator, who has very few resources to take, and who represents a vanishingly small percentage of the copyright violations going on. This doesn't mean that the enforcement won't be tried. However, the economics of this kind of law enforcement can already be seen in the war on drugs, as can its effectiveness. Different national jurisdictions just make this problem more difficult.

Finally, the measures needed to really prevent widespread copyright infringement basically involve building the legal and technical infrastructure for widespread censorship.

## 3 Alternative Funding for Copyrighted Works

In the previous section, we discussed why we don't believe that traditional copyright enforcement will work anymore, which means that many content creators will probably not get paid the way they traditionally have. This is likely to cause problems for many kinds of content providers, especially motion picture houses, since even a relatively cheap movie costs a great deal of money to make. (Novels can be written, and music written and performed, without a great deal of overhead. However, very few good movies can be made in someone's garage, and many very good movies simply could not be made on such budgets.)

If creators won't be paid through traditional copyright royalties, then it is worthwhile to consider what other funding sources are available. While we don't intend to make an exhaustive list, some alternatives are apparent:

1. *Voluntary contributions* Some people will commission works of art as they always have; some will be willing to donate money to see their favorite writer finish another book. The Street Performer protocol is one way this can be done.

2. *Advertisements* The content can be made available from servers that make their money from ads. If these servers are free, and are set up to do downloads of this content very quickly, then they may earn quite a bit of money, since most users will prefer getting the content for free from the fastest place, and won't mind seeing a few ads. Copyright enforcement is used, then, only against other sites that download content and try to resell it. We see this as among the most promising of the alternatives, and the Street Performer protocol can and should be used with it. We also note that many commercial web sites already do this, in some sense; e.g., the ads on various news sites and search engines pay for the availability of the sites. There are no mechanisms for preventing users from redistributing the content, other than occasional copyright warnings.

3. *Product placement* Product placement takes place when some advertiser convinces the content creator to make reference to his product or idea in some positive way. For example, many recent movies have had products placed in them as a way of earning additional income from the movie. We expect to see more of this. However, we note that it only works for some media, and that many creators and consumers will dislike such product placement, especially if it is blatant.

4. *Government funding* Many countries have some kind of government funding for the arts, and this may become even more common, if copyrights become very hard to enforce. However, there are obvious social consequences to having (for example) all novelists and musicians who are ever paid for their work paid by a government agency. It is also unlikely that even the most generous budget for funding these works will compare with the amount of money now spent on copyrighted books, music, movies, etc.

# 4 Our Solution: The Street Performer Protocol

## 4.1 Overview

Suppose an author wants to get paid for his next novel in an ongoing series. Using traditional commerce mechanisms, he would find a publisher who would effectively underwrite the creation of the novel. The publisher would then make the novel available to the mass market, in the hope that enough people will buy the novel to recoup his costs. If the author could not find a publisher, he could publish the book himself and hope to recoup his costs. In either case, the author and/or the publisher are taking a financial risk in the hope of making a profit.

There is an alternative. Using the logic of a street performer, the author goes directly to the readers *before* the book is published; perhaps, even, before the book is written. The author bypasses the publisher and makes a public statement on the order of:

> "When I get $100,000 in donations, I will release the next novel in this series."

Readers can go to the author's web site, see how much money has already been donated, and donate money to the cause of getting his novel out. Note that the author doesn't care who pays to get the next chapter out; nor does he care how many people read the book that didn't pay for it. He just cares that his $100,000 pot gets filled. When it does, he publishes the next book. In this case "publish" simply means "make available," not "bind and distribute through bookstores." The book is made available, free of charge, to everyone: those who paid for it and those who did not.

There are basically three things that can go wrong with this kind of system:

- The author can charge an inappropriate price. He and other authors will presumably learn from their early mistakes, and become pretty good at choosing appropriate prices.
- The author publishes the book before he gets the requested amount donated. This doesn't appear to hurt anyone directly except the author, but it may undermine participation in this kind of scheme later, especially in schemes run by this author later.
- The author gets his amount filled, but still doesn't publish the next book in the series. This will ruin his reputation for future deals of this kind, but that is only a concern if he has already built up a reputation, and if he intends to publish future books. It is here that we can see how to use cryptography and a trusted third party to make the whole system work.

The first two are marketplace issues, and essentially self-correcting. The third problem involves trust, and is one worth considering. The obvious way to solve this is to have a trusted third party handle the transaction. For lack of a better term, will call this third party the "publisher."

The author submits his novel, or parts of it if it a serial, to the publisher. The publisher has his editors review it to see if it's worth trying to sell (like any publisher, albeit with rather low printing/binding costs). If so, he and the author agree on a price and split. For unknown authors, the first several chapters, or even the first few books, may be freely available, in hopes of drawing in customers. For known authors, perhaps the first chapter or two is free, and the rest go through the payment mechanism. He has the whole novel, and on his web site, he makes available, say, chapters 1-3 for free, and chapter 4 will become available when $1000 is donated to the cause of getting it out, or on some target date.

Each donor of $N gets a signed certificate that's basically a kind of a security. On the target date, if this novel hasn't been released, then the security may be redeemed at the publisher's bank for $N plus interest.

The publisher can be as involved in the process as he wants. He could act as a traditional publisher, selecting, and editing, releasing, and promoting manuscripts. He would do this in the hope of extracting a higher price than the author could by himself, because of his publishing brand. He might also hope to make the novel appear first on his web site, and sell ads to make additional money. On the other hand, he could also be no more than a"vanity press," making no claims about the quality of the book and simply acting as an escrow agent for the author.

If enough readers want to see the next chapter, they can make a payment. The publisher needs no identification for this, so anonymous payment systems work quite well. The publisher holds the payments in escrow until the chapter is released, and then sends the author his cut.

Note that most of what is being done here is using a trusted third party to move the trust issues to some entity with a good reputation to maintain.

## 4.2 Motivation

The funding of the next novel in a series is a clear case of a public-good problem: each donor probably has very little impact on when or if the next novel is released. To understand some possible motivations, we must consider some situations in which street performers of various kinds get paid now.

1. A donor may give money partly out of the desire to be recognized as a nice person, a patron of the arts, etc.

2. There may be additional premiums involved in donating; raffles for a lunch with the author, for example.

3. A donor may be more likely to give money when he can see that it has an immediate effect. Thus, public radio stations have goals for pledge drives, and also for specific times. This might translate into letting novels out in dribbles, as small additional goals are met. Experience in the market will determine what pricing and marketing strategies work best.

# 5 The Street Performer Protocol

## 5.1 Roles

In the basic street performer protocol, there are essentially three parties: the Author, the Publisher, and the Reader (of course, the "Reader" is actually many people). Their aims are straightforward:

1. The Author:
   a. Wants to get paid the proper amount for his work.
   b. Doesn't want the Publisher to steal his work.
   c. Wants the publisher to adhere to any contracts, such as marketing and exclusivity.
2. The Publisher:
   a. Wants to get paid the proper amount for hosting the Author's work on his system, and administering the process.
   b. Wants the Author to adhere to any contracts, such as timeliness, exclusivity, etc.
3. The Reader/Donor:
   a. Wants the work to be published when sufficient donations are collected.
   b. Wants his particular donation to be reported properly, and for the author to get whatever percentage he and the Publisher have agreed to.
   c. Wants the "current balance" of donations to be reported properly.

Most of these goals are interpersonal, and can only be enforced by contract and the court system. Some, however, can be mitigated by the protocol.

## 5.2 The Protocol

Following is the basic flow of the Street Performer Protocol. We will assume that the work is a novel, and that it will be released chapter by chapter. We also assume that a Publisher is handling all of the financial transactions an will release the book. Of course, the same general protocol will work equally well for other types of digital property.

### 5.2.1 Submission of Work to Publisher

The Author submits some part of a work to the Publisher. This may include a whole novel, or just the first several chapters. She also provides the Publisher with the hash of the next few chapters to be published, and perhaps with the hash of all remaining chapters in the novel. The Author and Publisher negotiate terms, based on how much the next chapter (or several chapters) will cost to get released, and how the money collected will be split between the Author and the Publisher. When the negotiations are finished, the Publisher puts her first several chapters onto his website, along with a notice tracking how much money must be donated in order for her to release the next chapter.

Note that in some cases she will give the Publisher the whole novel; in other cases, she will give him only the first few chapters. It is even conceivable she won't be finished with the novel when she sells it to him, though this could put the Author and Publisher in a difficult situation, should the Author be unable to finish the novel in time. In the remainder of this section, we will assume the novel is written, and that the Publisher has, at any given time, the text for the next several chapters to be released. The hash of the final novel must be given to the Publisher at this point.

### 5.2.2 Gathering Donations

The Publisher gathers donations by, in some sense, taking bets on whether or not the novel will be released under various conditions. He sells donors a signed promise to return all donations, perhaps with interest, if the next chapter in the novel doesn't appear by a certain deadline.

The donor sends $X in donations, plus some unique identifier to specify where any refunds should go. Donors who wish to remain anonymous may specify either some anonymous account that goes back to them eventually, or some charity or other beneficiary of their choice. The only beneficiaries that should be discouraged are the Author and the Publisher.

The publisher sends a digitally signed document promising to repay the donation of $X, unless a certain event or set of events occur. The most obvious event to plan for is the next chapter failing to appear by the cutoff date. The next most obvious event is the last chapter in the book failing to appear by some longer-term cutoff date. The donor holds this signed statement, thus getting both a guarantee that he will be repaid if the Author refuses to release the work by the promised time, and proof that he donated $X for this work.

### 5.2.3 Paying Back the Donors

The donations are held in escrow until all conditions are fulfilled. Because the conditions are easy to understand and prove (they include a hash value of the material to be released), this can be objectively determined by just about anyone. If the promised work is not released by the specified date, then the donors' signed documents can be used to collect money from the Publisher. If he resists paying, the donors can ruin his reputation by showing that he didn't abide by the agreement.

### 5.2.4 Delivery

Once the required value of donations are received, the Publisher releases the chapter into the public domain. He could place it on his web site, and then inform the donors that the work is available. Ads on the site will presumably raise additional money.

## 5.3 Variations and Refinements

The basic goal in all on these refinements is, whenever there's a party with a financial conflict of interest, to replace him with someone who is paid a flat fee for carrying out a function, and isn't incentivized to conspire with any internal party.

### 5.3.1 The Banker

We can add a Banker to handle payments. We would have to modify the protocol so that he holds donations in trust for the Publisher. Bankers have a huge amount of reputation capital, and no financial incentive to cheat either the Readers or the Publisher. If payments work the right way, then Readers can send the Publisher their "receipts," which can then serve as enough proof to ruin the Banker's reputation if he cheats.

The Banker must not release the donation funds until the material is published. This must be precommitted to him: he's given the hash of the material to be published, the donations are accepted, he notifies the Publisher and Author when the desired level is reached, and when he sees it has been published, he pays up.

### 5.3.2 Story Content Manipulation

This covers a variety of items: short chapters, substandard chapters, requesting donations without having the content ready yet, etc.

All this stuff is handled by reputation. If the Publisher or the Author wants to build up or maintain a good reputation, then they must not do this sort of thing. Since the readers/donors will have direct recourse (stop donating money), this is enough.

# 6 Applications: Public Financing of Public-Domain Works

The Street Performer Protocol is effectively a means of collecting private financing for public works. It allows for all kinds of alternative public works: literary, music, video, etc. It can be used to improve public-domain software: companies could announce prices to add various features to an existing public-domain software package, and users could pay for the features they want; when enough people want a given feature, it gets designed and implemented. People could set up this protocol to pay for their web sites: if people are willing to contribute to a web site, then it will continue to be maintained and improved.

Another nice place for this is in terms of serials. People get really excited about television serials like Party of Five or ER, in which long-running ideas and stories are developed. It might be possible to keep a low-budget video series running for years by having a few episodes always queued up. The beauty of this is that advertisers and boycotts don't really mean much here: if enough people are willing to "vote with their (e-)wallets," then it doesn't matter how many angry Dan Quayle's supporters don't like Murphy Brown.[2] In effect, the United States Public Broadcasting System works in this way: people contribute money to see certain types of programming, but everyone benefits from what is eventually shown on the air.

# 7 Conclusion

The notion of an "author" who has "rights" to a "work" is a relatively new one, dating from the time of the printing press. Before then, it was impossible to separate a work from the physical instantiation of that work, so copyright had no meaning. Since then, the relative expense of copying and distributing works made copyrights possible, and led to their enforcement. Future technological developments will make copyrights unsustainable because the barrier to copying and distributing drops to zero. It will become impossible to talk about a physical instantiation of a work as something separate from the work itself because there can be arbitrarily many instantiations.

The Street Performer Protocol is obviously not a complete solution to the problem of marketing intellectual property in the age of free and perfect copying, but it is useful in some situations.

If a trusted intermediary administered the system, it could be implemented with no trust between the Author and the Reader. Authors who might have no publishing avenues in traditional media could release a sample of their work and solicit donations for "more of the same." In this way, the ability of the net to congregate people of similar interests could be used to finance works that might not otherwise be financed.

# 8 Dedication

This paper is dedicated to Ross Anderson, who spent some of his youth busking on the streets of Germany with his bagpipes.

# References

[And96a] R. Anderson, ed., *Information Hiding, First International Workshop Proceedings*, Springer-Verlag, 1996.

[And96b] R. Anderson,"The Eternity Service," *Pragocrypt '96, Part 1*, CTU Publishing House, 1996, pp. 242–252.

[BGH+95] M. Bellare, J.A. Garay, R. Hauser, A. Herzberg, H, Krawczyk, M. Steiner, G. Tsudik, and M. Waidner, "iKP - A Family of Secure Electronic Payment Protocols", *The First USENIX Workshop on Electronic Commerce,* USENIX Association, 1995, pp. 89–106.

[Bra93a] S. Brands, "Untraceable Off-line Cash in Wallets with Observers," *Advances in Cryptology—CRYPTO '93 Proceedings,* Springer-Verlag, 1994 pp. 302–318.

[Bra93b] S. Brands, "An Efficient Off-line Electronic Cash Systems Based on the Representation Problem," C.W.I. Technical Report CS-T9323, 1993.

[2] Of course, this works both ways; there's no doubt a market, albeit a small one, for a couple of KKK serials involving the touching story of a loveable bunch of goons burning crosses in the yards of people with the wrong skin color. Freedom of speech cuts both ways.

[CR93] CitiBank and S. S. Rosen, "Electronic-Monetary System," International Publication Number WO 93/10503; May 27 1993.

[Fer94] N. Ferguson, "Extensions of Single-Term Coins," *Advances in Cryptology—CRYPTO '93 Proceedings*, Springer-Verlag, 1994, pp. 292–301.

[FY92] M. Franklin and M. Yung, "Towards Provably Secure Efficient Electronic Cash," Columbia Univ. Dept of C.S. TR CUCS-018-92, April 24, 1992. (Also in Icalp-93, July 93, Lund Sweden, LNCS Springer-Verlag).

[LMP94] S. H. Low, N. F. Maxemchuk and S. Paul, "Anonymous Credit Cards," *The Second ACM Conference on Computer and Communications Security,* ACM Press, 1994, pp. 108–117.

[MN93] G. Medvinsky and B. C. Neuman, "Netcash: A Design for Practical Electronic Currency on the Internet," *The First ACM Conference on Computer and Communications Security,* ACM Press, 1993, pp. 102–106.

[NM95] B. C. Neuman and G. Medvinsky, "Requirements for Network Payment: The NetCheque$^{TM}$ Perspective," Compcon '95, pp. 32–36

[Oka95] T Okamoto, "An Efficient Divisible Electronic Cash Scheme," *Advances in Cryptology—CRYPTO '95 Proceedings*, Springer-Verlag, 1995, pp. 438–451.

[OO90] T. Okamoto and K. Ohta, "Disposable Zero-Knowledge Authentication and Their Applications to Untraceable Electronic Cash," *Advances in Cryptology—CRYPTO '89 Proceedings*, Springer-Verlag, 1990, pp. 481–496

[OO92] T. Okamoto and K. Ohta, "Universal Electronic Cash," *Advances in Cryptology—CRYPTO '91 Proceedings*, Springer-Verlag, 1992, pp. 324–337.

[ST95] M. Sirbu and J. D. Tygar, "NetBill: An Internet Commerce System Optimized for Network Delivered Services," Compcon '95, pp. 20–25.

[SK96] B. Schneier and J. Kelsey, "A Peer-to-Peer Software Metering System," *The Second USENIX Workshop on Electronic Commerce*, USENIX Association, 1996, pp. 279–286.

[TMSW95] J. M. Tenenbaum, C. Medich, A. M. Schiffman, and W. T. Wong, "CommerceNet: Spontaneous Electronic Commerce on the Internet," Compcon '95, pp. 38–43.