Unleashing the Power of Digital Goods: Enabling New Business Models for the Music Industry

Dimitri do B. DeFigueiredo

August 2003

1. Motivation and Background

Recent developments of peer-to-peer (P2P) file sharing applications coupled with advances in other technologies, such as data compression and high bandwidth communications, have altered the landscape of many markets where information goods are sold. The music and software industries seem to be changing the most rapidly, but many others (such as publishing and trading) are also affected. The technological changes may have significant economic effect. However, the trends that will set the path to be followed are not yet clear in many aspects. For example, one may ask the questions: How will the music industry make money in the future? Is the Free Software Movement going to maintain its present momentum? Despite its claims, it is not even clear that the music industry is being hurt by the unfettered copying of copyrighted songs that takes place through P2P networks[1] [28, 45]. Luckily, though there are many large uncertainties concerning these issues, careful analysis is able to provide some useful insight. Up to now, the music industry appears to be the most affected by the current technological changes. We will focus on it because it provides good motivation, its dynamics can be understood with relative ease and the discussion is timely; but we believe that most of the aspects discussed here are also applicable to other information goods.

Currently, the business model used by mainstream music corporations is based on the technological paradigm that was established a few decades ago when it was cheap for firms to reproduce vinyl records and tapes; but copyright (or some other means[2]) made it prohibitive for others to follow suit. Unfortunately, through some evolution of the market after the adoption of the present business model an oligopoly has formed. Five big firms: BMG, EMI, Sony, Universal and Warner have approximately 80% of the market [36]. These firms have established promotional mechanisms and

[1] Although CD sales declined by 6% in 2001 and a further 9% in 2002 [38], the average CD price has risen considerably in the same period (despite the present economic downturn) and the number of new releases has also declined [45]. Thus, one would be unwise to conclude with no further evidence that the decline in sales is due to P2P networks.

[2] Another factor maybe loss of quality due to successive copying.

distribution chains that seemed very efficient a few years back [36]. However, with the advent of P2P systems these methods of distribution and promotion do not seem as efficient from both the producer's and the consumer's point of view. In both cases, there are advantages to be obtained by using P2P networks. This new technological paradigm coupled with the large costs of copyright enforcement, may provide a brief window of opportunity that either allows for other agents to enter the market or for a change in the business model used. This text reviews a few of the new business models that have been proposed and explores the mechanisms needed to enable them.

It is interesting to note that record companies (or labels) are not themselves creative entities able to produce songs. Record companies rely on the creativity of artists to provide the goods consumers want. The artists on the other hand need the promotional campaigns and distribution chains provided by the companies. One could argue that record companies are successful at producing *stars*, artists whose work is considered worthy of recognition and reward. Once such reputable status is achieved (be it through a strong marketing campaign or other means) the artists acquire large bargaining power and are able to strike better deals with the record labels. In fact, once the artist becomes known it might be in his best interest to skip the record company altogether[3]. An unknown but successful newcomer is a much more profitable proposition for a record label than an already established and well known artist. The problem with that approach is the difficulty in finding such artists. Currently, the industry's success rate is very low [38]:*"Of thirty thousand CDs that the industry released last year* (2002) *in the United States, only four hundred and four sold more than a hundred thousand copies, while twenty-five thousand releases sold fewer than a thousand copies apiece. No one seems to be able to predict which those four hundred and four big sellers will be."*

As can be seen the underlying sales structure of the music market is very unevenly spread. The reasons for this uneven distribution have been studied previously [2], here we wish only to point out that better diversity (*i.e.* a wider choice) would be a desirable quality of new business models. Note that even though a CD that sells a thousand copies is a disastrous venture for a record company, artists may be happy to know that a thousand listeners enjoyed their work (specially if they can support it further). A thousand may be a significant number depending on how one looks at it.

The remaining sections are organized as follows. Section 2 surveys the related literature and provides a general understanding of the assumptions economists have made when modeling businesses that produce information goods. Section 3 describes new business models that have been proposed. Section 4 describes a few mechanisms that would be required in some form or another to implement the new business models. Section 5 describes one such mechanism in detail. Section 6 makes concluding remarks and proposes a few research directions.

2. Related Work

The literature on the economics of information goods is extensive[4]. A good reference is [40] where many aspects[5] of information goods are explored within a

[3] Something many do in an indirect way by creating their own record label.
[4] See [0] for a bibliography on the economics of copyright.
[5] Such as network externalities and lock-ins.

business mindset. A few comments may be made about the general treatment of the subject, as noted in [28]:"*Economists have tended to focus on the tradeoff between consumption efficiency (maximizing the amount consumers get of any intellectual product) and production efficiency (preserving incentives to create these products).*" [11, 24, 46].

From an economic point of view, since the coming of P2P networks goods that can be put into digital format have attained two interesting characteristics:

- **Non-rivalry**. An individual's consumption does not diminish the amount of goods still available[6].
- **Non-excludability**[7]. Once produced it is impossible to prevent anyone from consuming the goods available.

These characteristics are precisely the characteristics of so called *public goods* (*e.g.* public lighting or national defense) [44]. Thus, one may analyze information goods as public goods. Until recently it was thought that private provision for public goods was inefficient, in the sense that not enough of those goods would be provided even though everyone would benefit. Fortunately, [9] showed that this is not necessarily the case. Thus, it is possible to have an efficient allocation of private resources to produce those goods. However, they do not specify a mechanism through which this could happen.

A common metric to analyze the efficiency of the different business models is social welfare. In general, social welfare is defined simply to be the sum of total consumer surplus (consumer utility minus the price of a product for all consumers) and gross producer's profit. Whether giving Sony a million dollars is just as good as giving one extra dollar to 1 million of its customers is debatable, but this is the general framework within which many authors attempt to compare the different possibilities [11, 24, 46].

Another common assumption in the literature is that copies have a lower value than originals [32]. This assumption is generally made to provide consumers with some incentive to buy originals instead of just copying them for free [11]. Although, this may not always be the case, it is unclear how else to provide such an incentive for consumers to reward producers and to model it in economic terms. However, this assumption is restrictive precisely because it views copies as lower quality content. This may simply not be the case. Some papers have been able to provide very new ideas and insights through the use of non-standard assumptions. For example, if copies and originals have the same value, a gift economy[8] may provide the incentives to reward producers. Such a model is a considerable departure from the current perspective and a tint of skepticism is adequate. Yet, one should not simply dismiss any efforts in these directions as they do bring very new ideas and insights. As an example, Grothoff [21] was able to provide a very simple framework for resource allocation in a P2P network through the use of such assumptions (which in that particular case may be quite appropriate).

The literature generally assumes that producers are rewarded exclusively

[6] This was also true before P2P. If I have a CD and a friend of mine makes a copy of it, I may still listen to it whenever I want to and the same is true of my friend. Obviously, depending on which CD we copy this may be illegal.

[7] This might be the principal change that has occurred to digital goods with the onset of the Internet [24].

[8] See [30] for a characterization of a gift economy.

through the purchase of originals. There is a little twist in the music industry that may make this assumption inapplicable: artists also earn money through other means such as concerts and marketing licenses[9]. In some cases, the rewards obtained through these separate channels are considerable and should be taken into account. While the loss in revenue due to illegal copying represents decreased income, the extra publicity obtained through illegal copying might increase the revenues obtained through other means.

Within the standard framework, the work of Varian [43] provides some interesting insights at the possible alternatives available to a profit maximizing producer. It is interesting to note that the framework used by Varian assumes the producer is able to exclude consumers from using the goods produced. This may be true for Video Rentals, but cannot be regarded as a realistic assumption for music at the present time. The results obtained are therefore either applicable to information goods other than music or can be seen as what could be achieved through a DRM[10] scheme able to exclude non-paying users from enjoying the content at hand. The paper analyzes the possibilities from the (profit maximizing) producer's point of view and pays no attention to social welfare implications. However, clear guidelines for profit maximization are established[11].

The social welfare implications of different business models are the subject of [11, 46] who independently developed very similar models. Belleflamme's analysis follows a path of increasing complexity by increasing the number of producers and finds different results as complexity increases. The results are similar to those obtained by Yoon [46] for the simplest models and show an increase in social welfare at the expense of a decrease in the producer's profit[12]. Yoon goes one step further in the discussion of social welfare and establishes within his framework the optimal level of copyright protection from the producer's point of view and from a social welfare point of view. The level of copyright protection is equated to the consumer's reproduction cost and fixed for all consumers. The two papers have inconclusive findings for more elaborate constructions. Belleflamme is unable to determine an equilibrium point and Yoon shows (within constraints) that the desirable level of copyright protection depends on many other environmental factors.

The case of P2P networks is studied with some detail in [17]. The study models the behavior of two firms. One firm distributes its music through the traditional business model incurring a high fixed initial cost for marketing, promoting, distributing and inventory. The other uses the new P2P networks to distribute its content. One notable assumption made in the paper is that the new distribution technology off loads the distribution cost onto the consumers themselves. It is the consumers who search and download the desired material. The firm using P2P technology does not have to pay such high fixed costs and is thus able to offer artists with a smaller audience (who generate smaller revenue and would not be profitable in the traditional model) with a feasible distribution channel. This increases the overall diversity of products available. Within the model and some limitations, [17] does find that the incentive to create is more efficiently employed when there is copyright infringement and thus copyright may be loosing its original purpose of encouraging new creation. The findings are in partial agreement with

[9] Yo-yo Ma and Britney Spears are good examples respectively.
[10] Digital Rights Management (or Digital Restrictions Management)
[11] A good bibliography on many aspects of DRM can be found at www.giantstepsmts.com/drmbiblio.htm.
[12] The benchmark comparison is against a monopolistic producer.

[46] and [11] and state that producers using the traditional model have a decrease in profit. However, this decrease is offset by a large increase in consumer surplus which in turn increases overall social welfare. The analysis of [17] compares the strategy of the two firms, one using P2P for distribution and the other using traditional means.

Observing that once the first copy of an information product is produced it is very hard to prevent further copying, [27] suggests a new business model called the *Street Performer Protocol*. Individuals contribute towards a pre-set goal that if achieved, allows the producer to make the product available. The idea has some very desirable properties: it allows for price discrimination between consumers, it eliminates copyright enforcement costs and it also maximizes consumption efficiency by allowing anyone to use the product once it has been produced. This model will be examined in more detail in the next section. In a similar setting, the work of [19] examines auctions for digital goods and concludes that some auctions fair reasonably well. In most cases however, the performance from the producer's point of view is worse than that of optimal fixed pricing. Also assuming the enforcement of copyright, Fishburn and Odlyzko [18] establish an optimal pricing strategy for competing firms selling digital goods. The results suggest that bundling and payment through a fixed price are the best options.

The extremely uneven distribution of wealth amongst artists is the subject of [2]. The text argues that this uneven distribution is due to two market characteristics:

- The author assumes that *"appreciation increases with knowledge"* and models exposure to art and discussion about art as consumption. Thus, the more a consumer can discuss about a song he likes with his friends the more value that particular song attains. (This leads to everyone enjoying *the same* few artists – the "stars".)
- The reproduction costs of the goods do not scale as fast as the potential market. In other words, there should be a small marginal cost.

Adler's work is successful in showing that there can be huge differences in income for individuals with identical "talent"[13]. The work separates "talent" and income. Adler points out that factors other than talent are going to decide who will be a "star". These findings are consonant with reasoning presented in [36] that attributes to promotional and marketing muscle the greater part of the success of artists with little reputation.

There is also considerable literature about the implementation mechanisms mentioned in latter sections. Electronic payment systems, for example, have been the subject of intense research for sometime. A good semi-technical overview of the systems that have been proposed can still be found in [6]. The work of Chaum on digital cash [15] also deserves special attention. Chaum is able to establish anonymity for purchasers just like what would be obtained with real cash[14] as long as individuals do not spend the same electronic cash more than once. Quite a few schemes have also been proposed to deal with payments of very small economic value, *i.e. micropayments* [23], the need for such systems occurs due to economic constraints. Considering the present alternative of credit cards Lipton and Ostrovsky [29] mention that: "The cost per such transaction is about 10 cents, and hence is not financially viable for tiny-cost transactions. Moreover, since the Bank must maintain 99.99% availability, even during peak traffic time (…) this requires additional cost in order to maintain capability for additional throughput and backup

[13] In fact, an individual with less "talent" can earn more.
[14] A good bibliography on Chaum's work can be found at [14]

systems." Thus there is a need for micropayment systems, if payments of smaller monetary value are to be made.

The proposed micropayment systems have focused on the trade-off between complexity of implementation (and hence transaction cost) and fraud prevention/detection with a bias towards the former. A few very ingenious probabilistic schemes that are able to aggregate payments have been proposed [25, 31] and even implemented to some degree [35]. However attractive they may seem, the proposed schemes have undergone considerable criticism either individually [41] or as class of payment systems as a whole [34] and have failed to be adopted by the public at any scale. The *de facto* standard payment system used on the Internet today is credit card transactions[15]. Despite their inherent lack of security, their convenience and widespread adoption seems to imply that they will be around for a very long time [10]. Thus, any new payment system should be feasible and competitive taking into account the prominence of credit card payments.

3. New Business Models

Between the many possible business models available to a producer of digital goods, we chose to focus our attention on those that do not rely on excludability. We believe that at the present time it is unfeasible (or very expensive) to exclude non-paying customers from enjoying information goods. Thus, we focus on business models that use non-excludability and the Internet to their advantage as opposed to seeing it as a threat.

Among the first possibilities that spring to mind is 3rd party sponsorship. Business models where producers are paid by a 3rd party have been around for a very long time. Open TV and radio are good examples of businesses completely based on such sources. It is interesting to note that such models are intrinsically well suited to the non-excludability characteristic we take to be prevalent in the current scenario. It is usually in the best interest of a sponsor that a sponsored artist's work be as widely distributed as possible even when the sponsor is targeting a specific market niche. However, we also believe that this model has been adopted by internet-based businesses to exhaustion; the number of businesses vying for those dollars is large and range from search engines to free web hosting. Thus, it is probably not wise to focus exclusively on this source at the present time. But, this source may be sought to complement other means of income especially if especially favorable circumstances arise[16]. Furthermore, implementing the models described here – which obtain funds from the consumers themselves – does not exclude the existence of sponsorship.

Both models described next assume that non-excludable goods are produced. Both also provide efficient mechanisms against under-consumption. Surprisingly, under the (albeit unrealistic) assumption of perfect information the SPP model can also generate more profit than optimal fixed pricing in a perfectly monopolistic market (due to the ability of the producer to price discriminate between the consumers). To provide this tolerance to non-excludability different assumptions and

[15] According to [12] as recently as 2002 over 80% of all non-recurring e-commerce payments were done using a debit/credit card.

[16] This assessment is circumstantial and may change significantly in the future.

compromises are made in the two models. There are, of course, may other possibilities.

3.1 The Street Performer Protocol Model

Through an analysis somewhat similar to that followed in [9] Hougaard *et al.* [24] showed that allocation of resources through the Street Performer Protocol (SPP) can be very efficient. "*In other words, within the framework of the model, selling information/digital goods does not constitute a problem with respect to profitability and efficiency - it only calls for a different selling mechanism.*"[24] This potential efficiency against both underproduction and underutilization may be unique to a class of business models similar to the one proposed by [27].

The proposed business model is simple. "*A provider (or vendor) announces a fixed total price for some information good and a date by which the price should be paid. Each potential user decides how much to contribute. If the sum of their contributions exceeds the total price then the information good is made available on the Web – otherwise not and contributions are returned.*" [24] Clearly, not all products are suitable for this model. For example, with no further refinements it would be unfeasible to sell news stories through this business model. One of the assumptions made here is that users will behave as investors committing resources to a product they will only be able to obtain in the future. This may require a significant change to a consumer's mindset or some form of credit provided by yet another agent[17].

A detailed analysis of this business model is provided in [24]. Clearly, the producer needs to have a good estimate of the total willingness to pay for the product in order to maximize its profits. This task may be significantly harder than just finding the optimal fixed price. To show that the business model proposed is Pareto efficient, the analysis presented in [24] assumes all agents know about the willingness to pay of all other agents. Needless to say, this is usually not case. However, mechanisms may be put in place to increase the amount of information available to the agents. In fact, the analysis in [24] suggests that any mechanism that will increase the valuation information available to consumers will actually make the producer more profitable.

For some goods we can solve the problem of having customers behave as investors by letting them experience the goods before they are actually called upon to finance them. This basically implies that the customers develop a "sunk cost" mentality which will encourage consumers to contribute by having the threat of not ever being able to enjoy the complete works while at the same time providing the benefit of letting users know their own valuations for the goods to be produced. We can draw a parallel with the case of the Blender software package [33] to be presented later. There, the money gathering efforts may have been very successful due to the threat of losing "sunk cost" in the software. Unfortunately, this may only work for a few types of goods. It is unlikely this would work for music for example. Similarly, the introduction of movie trailers or alpha software releases will enable consumers to find out about their own willingness to pay. If this is also followed by the introduction of user groups and fan-clubs respectively, users may be able to determine the willingness to pay of other agents. Both of these mechanisms tend to discourage free riding and make the producer more profitable.

[17] Note that there is an interesting twist to the credit alternative as credit is being provided for contributions that may ultimately be returned.

Free riding may be the most significant problem in business models that maintain non-excludability. Further discouraging is that fact that without the complete information assumption, in SPP it is rational not to contribute [24]. Each user wants to have the goods produced while at the same time contributing the least. With a fixed deadline, a possible strategy is to hold-off contributing for a while and wait to see if the total contribution of other parties will suffice to have to goods produced. A general method for avoiding such pitfalls is to make every potential consumer pivotal to the production of the goods. In other words, if each consumer is aware that him/her not contributing may mean that the goods will not be provided, there is a strong incentive to contribute[18]. The Rational Street Performer Protocol described later tries to address this.

Two other related difficulties that must be tackled when using SPP are: the need for the threat of not producing the good to be credible and the certainty of producing the good if the goal is reached. If the threat of not producing the good is seen as a bluff, consumers may be tempted to hold off contributions in the hope that the same good will be offered for a smaller price at a later date. Because the producers of information goods are likely to be in business for longer than just once, this may cause some problems. Also, consumers should be confident that if the target is reached the goods will be delivered. Otherwise, the value of the goods may decrease significantly.

A good example is the case of books. Imagine an author is willing to use SPP to fund his work. The author also chooses to let readers read a few chapters before they have to decide whether or not they want to contribute. Readers need assurance that if the target amount is reached a whole book will be available otherwise they may be unwilling to contribute. It would virtually useless to have just a few chapters more. Thus, despite the considerable undertaking, an author would need to commit to producing a whole book once the target is reached. If an author has been unable to fund his last few projects, it would be tempting to offer an already written title for a cheaper price in order to make ends meat. Thus, it would be better for the author to sell his books chapter by chapter and not have the "sunk cost" of developing a whole book.

A few extensions of the Street Performer Protocol have been proposed. The *Rational* Street Performer Protocol tries to make it rational for consumers to contribute. The protocol is very similar to SPP except that contributions occur in a series of rounds and pledges are a function of the total amount raised. In other words, contributors make pledges of the form: "I will donate one dollar in every $____ raised over $____ up to a maximum contribution of $____" and after each round "The maximum amount that could be collected while abiding by each contributor's requirement is calculated." Money is collected from the contributors and the good is produced if the goal is reached or given back if it is not enough [22]. The protocol hinges on the idea that each person's pledge is a function of the total amount raised. The author of RSPP claims that in this variation of SPP, it becomes rational (after a few rounds) to donate money to "the production of the public work". However, the analysis assumes that an individual's utility function of the total amount of money raised is a monotonically increasing function. Thus, the analysis presented does not apply to our case where the utility of the total amount raised is a monotonically non-decreasing step function as the good is either produced or not.

Another variation of SPP is the *Wall* Street Performer Protocol which

[18] As noted in [24], this is in effect what complete information does to make the Street Performer Protocol Pareto efficient.

attempts to refine SPP to enable funding of open source software development through a bond market. Instead of allowing for producers to determine what can be produced, the protocol allows for consumers to determine what they want produced by issuing bonds that are paid off once the good is made available. In this proposal it is the producers – not the consumers – that have to gather to provide the goods; the payment has already been set aside.

Another possibility (which we have not seen in the literature) is to have a variation of SPP where the time the good produced is made available depends on the amount of money raised. For example, if a total of one million dollars is raised the good becomes available in 2 weeks. If only $500,000 is raised then consumers need to wait 2 months. Finally, if only $100,000 is raised the waiting time increases to 2 years. Furthermore there may be different deadlines for different money targets and therefore the money gathering effort need not stop when the first target is reached. It might be feasible to speed up production or delivery if a higher target is reached at a later date. This modification provides a credible threat of not producing the goods by delaying the release considerably. For example, if an author has already written a book but the target is never reached it would be unsatisfactory to *never* release the book as would be required to provide a credible of not releasing future works either. This would be wasteful as only a few of the total number of creative works produced would ever be enjoyed. With the delaying variation the book may still be released, but in 20 years.

A very interesting case study of a real world implementation of the Street Performer Protocol is given in [33]. There the case of the 3D freeware Blender is considered in detail. The software package had been developed by a company that went bankrupt and thus there was a threat that the software package may become obsolete with time. The copyright holders set a price of 100,000 euros to release the software as open source under the GNU GPL. Under a newly formed foundation, organizers were able to raise the required amount and the software as released under the GPL in a surprisingly short length of time.

3.2 The Distributed Patronage Model

Another model is perhaps best described through an example. Suppose a musician, Mary, records a CD that will be available at the shops for $15. Two of her potential customers are Alice and Bob. However, Alice and Bob have some very different background. Alice appreciates Mary's work and is doing very well financially, she would be willing to pay $20 for the CD. On the other hand Bob's earnings have been severely hit by the dot-com bubble and he is only willing to pay $3. With the current business model, Alice will purchase the CD for $15 and thus keep $5 while Bob will go home unable to listen to Mary's work[19]. It is interesting to note that if the cost of copying a CD is negligible, it would not have cost Mary anything to keep Bob's $3 and give him a copy. This would be good for both Mary and Bob.

From the producer's point of view and to maximize efficiency against underproduction, it would be good to have the ideal situation where each party gives their utility values to "purchase" the digital goods, independently of how others value them.

[19] Social welfare would be $p + 5$, where p is Mary's profit and 5 is the difference between Alice's utility value and the price paid (20-15 = $5).

We would thus have a system in which the same CD sells for different prices depending on who buys it. The following figure illustrates the two different scenarios.

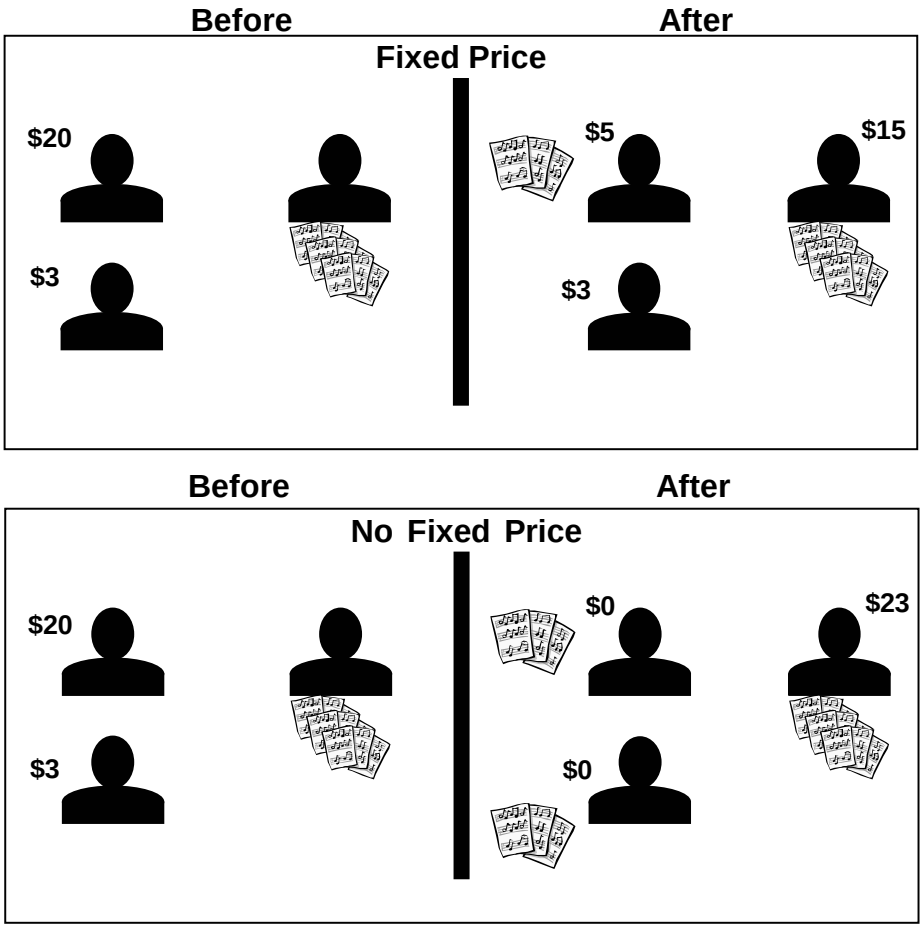

Fixed versus Variable Pricing

There is a very well known transaction in which the more efficient variable pricing mechanism is common practice even though all the parties involved have complete information about each others transactions. In our example, this means that Alice would know that Bob only "paid" $3 for the same good that she "paid" $20 for. However, she would not have any restrictions to that. Similarly, Bob and Mary would not find inadequate that Alice paid more than Bob for the same CD. *This variable price complete information scenario occurs when donations are made.* When donations are made to the producers of digital goods we achieve this "ideal" variable pricing mechanism. In our example, this translates to Mary distributing her music over the Internet and *suggesting* that people give her donations for her to continue her work. Note that, in this scenario, it is *legal* for anyone to freely copy and redistribute Mary's work. Furthermore, it is in Mary's interests that as many people as possible copy her work. Everyone is able to enjoy Mary's music even if they do not have any money to give. It is clear that, abundance is a key element of this system; Mary's work becomes available in unlimited supply.

Clearly, because economic agents are bound to provide less money for donations than they would for payments, this business model may be at a severe disadvantage when compared to the conventional model. That is an intrinsic disadvantage of this new business model from the point of view of the producer if we consider copies "sold" as the only source of income. However, in the specific case of the music industry this can be offset by two factors. One is the increase in revenue from other sources that is generated by the extra publicity. Another is the large amount of money that goes to the middleman (*i.e.* producers, lawyers, managers and the oligopolistic record companies). In the case of music, the average starting musician would be hard pressed to get 10% of the revenues from CD sales as royalty[20]. Going back to the example, even if Alice and Bob are not willing to donate $20 and $3 dollars respectively for Mary's work; a *tenth* share of that, respectively $2 and 30 cents would give Mary *the same* financial return as long as those quantities do not have to be shared with other third parties.

The example shows that this business model may be of great benefit to Mary, Alice and Bob as they are all better off by using it. Mary gets $2.3 instead of $1.5 (10% of $15), Alice pays $2 instead of $15 and Bob pays 30 cents but is now able to enjoy the work. This gain is obtained through eliminating inefficiencies (here represented by middleman) in promoting and distributing music. In other words, the great final outcome obtained by using the new business model in the example stems from an underlying assumption that promotion and distribution can be done through simply copying files in a P2P network. Obviously, P2P networks have been an enormously successful means of *distribution* as can be seen by the success of P2P networks such as Napster and Kazaa. However, it is easy to download a song only once the song attains publicity and one knows what to look for; not before then. It is not necessarily easy to promote a song on the Internet.

The model presented here is very similar to the already established idea of freeware in the software world. Freeware is software that is freely available for anyone to use. However, users are encouraged to send payments to the authors; usually by sending a check to an address that is specified in a text file shipped with the program. It is believed that if the users find the program useful they should send money because it is "the right thing to do" and also to maintain the software so that it can evolve with time. Freeware benefited greatly from the introduction of toll-free numbers for credit card payments [3].

Summarizing, the distributed patronage model provides for very good consumption efficiency as all consumers are able to enjoy any song. However, for it to become feasible and provide enough incentive for artists to "produce" songs; it is necessary to have mechanisms that *promote* and *distribute* songs and also enable consumers to make a large number of small donations through the Internet. These enabling mechanisms are the subject of the following sections.

4. Mechanisms

Now that we have established what business models we want to encourage, we need to answer a few questions: What is required to enable these new business

[20] The contracts signed by starting musicians tend to be very advantageous for the record labels, one of the reasons for that being that a starting musician has no bargaining power. See [8] for a description of a "standard contract".

models? What is already available on the Internet? Can artists use the Distributed Patronage Model or the Street Performer Protocol to fund their work at the present time?

Unfortunately, at the present time it would be very difficult for an artist to fund his work through the Internet. For example, consider the situations that would arise when using the Distributed Patronage Model. Even a resolute fan would find it very inconvenient, if not impossible, to search for an address and then send a check through the mail. However, with a little help it may become a lot easier to contribute. If the artist sets up a simple webpage with a Paypal[21] "donate" button and embeds a link to this URL in his freely available MP3 file, anyone with a Paypal account now has a convenient way to donate money. To be more precise, not everyone but only those who:

- Have a Paypal account,
- Can read the embedded information,
- Have internet access available at the time they want to make the donation
- Are sure no one else changed the link in order to get the money.

All they need to do is:

1. Read the embedded URL.
2. Open their web browser and go to the desired address.
3. Decide how much to give.
4. Click on the donate button and type in their password.
5. Wait for the confirmation screen.

Although the final procedure may take only a couple of minutes, by enumerating the tasks required it is easy to see that it is still a lot more convenient **not** to contribute.

As can be seen in the example above, a few tools are required to help both producers and consumers to implement the business models described previously. A payment system stands out as the most immediate requirement, but other mechanisms are needed also. These mechanisms may implement tasks that could not be performed previously or may simply make other tasks more convenient.

Because there appears to be no clear cut solution and each business model could be implemented in a variety of ways, in order to obtain a list of mechanisms that would be required to implement the new business models we considered how those business models should work in real life and then worked our way back to the present situation. We will follow that analysis here.

4.1 Promotion

Assume Alice wants to get some music. Until a few years back she had to go to the music store, so let us stick to this paradigm for a minute. If Alice decides to get new music it is either because she heard something she liked, and decided to get it now; or that she feels the need to hear something new and is going to the store to find it. In one case, Alice knows exactly what she wants in the other she does not. We can only guess which of the two possibilities happens more often. But there is a clear distinction here between Alice's experiences prior to going to the music store. In one case, she has been exposed to some "promotional" material and in the other she may not have been. It is important to note that being exposed to content is a powerful instrument towards getting

[21] Amazon's honor system would work just as well.

someone to like music; hence the importance of promotional campaigns. We conjecture that even though both possibilities occur when downloading music from a P2P network, users of P2P networks tend to look for material they have already been exposed to[22]. Although P2P networks are improving continuously, their evolution tends to address mostly the distribution problem, there are very limited means to promote music on the Internet. In fact, those tend to mimic the promotional methods that existed before the Internet and be very inefficient. For example, to cater for its older audience a pop radio station may want to play some old rock despite the fact that many young listeners are going to change the dial. Internet radios are able to provide each listener with a customized selection of songs, but not all do and thus suffer from the same problem. In the great variety of conventional promotional channels available: the radio, TV commercials, suggestions from a friend, opening a concert for another artist, etc; we note that those channels that are targeted at a single individual are usually the most persuasive (*e.g.* a suggestion from a friend) precisely because they are usually the most reliable.

New promotional mechanisms that give artists exposure and are able to promote someone with little reputation are needed. The internet provides the ability to customize promotional campaigns to a single individual, something that was not possible previously. This may drastically increase the efficiency of promotional mechanisms but also raises new privacy issues as one needs to gather data about an individual to customize his experience. Customized Internet radios may be used for such purposes, but they do not harness the possibilities available through social networks. A P2P network that: knows your musical preferences, enables you to share your thoughts about songs you like, and is able to suggest new songs would be an interesting possibility.

4.2 Obtaining Music in the Future

Let us follow another hypothetic chain of events to try and observe what mechanisms are required by the Distributed Patronage Model. It is the year 2010. Charlie has just received an instant message from Bob, who is raving about this new song he heard on his customized radio station. The message arrived at Charlie's presence enabled cellular phone but he is with a group of friends and cannot hear it right now. However, an attachment to the message contains a permanent link that enables Charlie to be sure he will be able to get the same song later. When Charlie gets home he picks up his wireless notebook and opens up the instant message. He clicks on the link and automatically the song is downloaded and starts to play. He is immediately able to rate the song and/or send a message back to Bob saying what he thought of it. In fact, he can even find out when the band is playing in his hometown and download all other songs from the same band. While Charlie is enjoy his new song, behind the scenes (doing it with Charlie's consent) the MP3 player is using the Distributed Patronage Model and automatically sending small donations to the band. This is done automatically so that Charlie does not have to remember to do it all the time. Those small donations from many different listeners are also being added up and at the end of the week the results are posted in an electronic donations billboard. The band Charlie is listening to is now the number one donation receiver for Pop/Rock in Charlie's town. There is no danger that Charlie might

[22] Whence the standard “searching” interface available in many P2P systems.

spend too much money on donations. He set up his IM[23] and P2P enabled MP3 player to cap those at a maximum of $10/month. However, he can also click on the: "I love it!" ($0.25) or the "Very nice!" ($0.10) buttons as much as he wants. The controlling software can either keep giving out the donations until the total is reached and then simply ignore the rest, or it can (proactively) estimate how much Charlie will likely give out based on previous months and only send out a fraction that is proportional to the $10/month maximum. The software can also avoid sending the donations one by one and collect all requests; sending the aggregate values at the end of the month after making all the appropriate calculations (the only problem about that is that we only find out that Charlie liked the band so much a month later, too late for hitting #1 on the electronic donation charts!).

The small story above illustrates the large number of functionalities that could be provided in order to transform the internet into an efficient promotion and distribution vehicle for music. To see that we can enumerate a few of the functionalities mentioned:

- a permanent resource locator that links to a song;
- the ability to promote songs through "word-of-mouse";
- a payment mechanism that enables a large number of small donations;
- the ability to rate songs and get customized content;
- having a (customized or otherwise) billboard to encourage donations and
- a personal software agent that makes donations for you.

These are all tools that would encourage the adoption of the Distributed Patronage Model for the hypothetical situation presented. Had a slightly different model been proposed, other needs would have been found. For example, thus far it has been assumed that support is given to artists through their songs. Another possibility (which is closer to the traditional patronage model used extensively in the past) is to simply decide which artists one is going to support and then support those artists in the long run independently of the work they are producing at any one moment. This could be done by giving such artists a monthly donation and it eliminates the need to determine what songs are being enjoyed. Similarly, other tools and mechanisms would have presented themselves had the model discussed been the Street Performer Protocol.

4.3 Payment systems

The use of credit and debit cards to pay for music suffers from the fact that a significant number of listeners are of an early age or due to some other constraint are prevented from having such cards. Better provisions to accommodate this group of individuals may make a significant difference (a lesson quickly learned by the providers of prepaid cellular phones).

A payment system intended to be used by the Distributed Patronage Model has a very peculiar set of requirements as payments are (ideally) administered by untrustworthy autonomous agents and are considered "donations". To protect users and encourage use of the system, those payments (or donations) should be revocable. In fact,

[23] Instant Messaging

revocability is already present on the Amazon Honor System [5][24]. Furthermore, depending on the operation of the paying software agents the values of the donations made can be very small and the donations made by a particular agent repetitive. Micropayment schemes that aggregate repetitive payments such as Payword [37] may be very useful. In a further refinement, if one also wants to implement an electronic billboard to record the total amount received from donations by any particular artist, there should be a publicly verifiable mechanism that either aggregates the donations or receipts from those.

A payment system for SPP has a very different set of requirements from those of the Distributed Patronage Model. SPP works best if it can aggregate large amounts of money. Perhaps the only non-standard requirement in this case is that contributions be returned to contributors if the target amount is not reached. SSP requires the existence of a trusted third party[25] to perform the aggregation of all donations. Different cryptographic mechanisms such as secret sharing [39] may be used to simulate the trusted party. In particular, the need to achieve a credible threat of not releasing works if the target money amount is not reached can be achieved if the author is himself made unable to obtain a copy of the completed works. Encrypting the final version with the public key of the trusted third party and destroying all other copies would be a simple way to achieve such assurance. Clearly, more elaborate mechanisms may be useful (*e.g.*, zero knowledge proofs [20], optimistic fair exchange [7], verifiable encryption [13]).

A minimal requirement for implementing SPP would be a website with a "Donate" button and a running counter for the amount of money donated so far. But this neglects the potential we have to make more information available to the potential consumers and thus make the assumption of complete information more pertinent. As the analysis by [24] suggests: the more information available to the consumers the more profitable the producer[26].

4.4 Available Mechanisms

Perhaps the most prominent tools that are present in the Internet today and used to distribute music at a very low cost are P2P networks. The technology behind P2P networks is a field of intense research efforts. We refer the reader to [1] for an overview of P2P networks. Most of the desired characteristics of P2P networking that are important to the distribution of music are the same as those desired for any other content: availability, scalability, authenticity of content, reliability and security all play important roles. Recent implementations of P2P networks have started to explicitly address many of these concerns. For example, the KaZaa network uses a "participation level" system to encourage users to share files and thus prevent free-riding [26]. The network also has an integrity rating system for each of the files shared and some of the files distributed are guaranteed to be authentic and managed by a DRM system [4].

---

[24] Unfortunately, there is a lower limit of $1 for all donations and a substantial 15% charged by Amazon.

[25] Trusted parties have many other uses: to sign certificates of claim of authorship, to take part in a fair exchange protocol or to run a billboard with vying artists for example.

[26] An early attempt to promote SPP as a business model was made by openculture.org, they even achieved charity status (that made donations tax-deductible). However, due to legal requirements (on charities) they are unable to return donations.

Authors of digital works already have an easy way to choose what kind of licensing terms to use for their works. A range of possible licensing agreements has been made that spans the current (very restrictive) "all rights reserved" model to the completely open "no rights reserved". There is even an easy way to choose amongst the licenses online without having to read them all [16]. Through intermediate licensing terms musicians are able to provide consumers with free songs while still requiring – for example – producers of TV commercials to pay. Also available is a unique identification numbering system for recordings that encompasses both audio recordings and video recordings of audio content (such as concerts), the ISRC (International Standard Recording Code) has been in use for many years. These can be used in conjunction with ID3v2 tags (a metadata embedding mechanism for MP3 files that has also gained wide acceptance) and online databases such as the ones available at: musicbrainz.org, freedb.org or gracenote.com to obtain detailed information about content[27].

5. MagicThanks

Previously, it was noted that there is strong motivation to provide economic models in which it is legal for anyone to freely use, copy and redistribute digital goods. One such model is the Distributed Patronage Model, for it to become a viable alternative to the current system, a way for the users of digital goods to transfer money to the producers of the same goods is needed. There is no widespread system right now that is adequate to this task. A convenient electronic payment system capable of performing a large number of small donations is necessary for the Distributed Patronage Model to become viable. Also, it should be possible to locate authors of digital goods. This section describes **MagicThanks** a simple protocol that links authors to their works and enables the transfer of funds through the use of the existing Internet infrastructure. The system is proposed as a mechanism for the implementation of the Distributed Patronage Model[28].

5.1 System Description

For the sake of simplicity and efficiency we strive to achieve only very limited gains by using MagicThanks. The system provides a mechanism to locate authors of digital works. This is an important step in enabling the transfer of funds between authors and users. The actual payment mechanism used is outside the domain of the system. Furthermore, throughout this text we will assume transfers of funds to be donations and so there is no underlying concern to ensure finality of payment. A loose interpretation in how such transfers are to be provided will be taken. Such issues are thought to be implementation and application dependent just as whether it might desirable to make a payment revocable until its completion. The system is concerned primarily with giving money to the correct recipients.

The protocol does not directly prevent malicious parties from claiming authorship of works they did not create as in other DRM proposals. Thus, it is possible to profit from maliciously copying someone else's work and claiming to have created it. However, the system is able to detect such fraud and thus provides a strong incentive

[27] Interestingly, gracenote.com even runs a "digitaltopten" billboard system that is updated weekly.
[28] A somewhat simpler proposal (developed independently) is available online [42].

against it.

By using the proposed system, anyone who can send money to an individual's *contact address* can transfer funds. Currently, Paypal provides the ability for anyone to send money to e-mail addresses. We will focus on using this functionality, due to its large user base and present availability. But the system is able – through different implementations – to transfer funds between entities uniquely identified through means other than e-mail addresses[29]. Most of the functionality needed by MagicThanks should be provided by a client software agent, which can search the Internet for specific content and controls contributions to authors. These clients could be seen as the peers in a peer-to-peer network. To provide fraud prevention and detection capabilities the system assumes the existence of a reliable on-line time stamping system[30] and an identity certification authority for e-mail[31].

In the remainder of this section we provide a brief description of the system through the use of an example. In what is described bellow, all *contact addresses* are assumed to be e-mail addresses. A quick analysis of the system follows the description.

5.2 The Publishing Process

Let us step through the way an artist can use the proposed system for the distribution of her creative works. Suppose Alice is a musician. She has just finished the recordings for her first song titled: "My First Soul" and she wishes to make the song available through the Internet. Alice also wants to be able to receive donations to continue her work. How should she proceed?

First of all, Alice needs to obtain an e-mail address[32] and a certificate asserting that she is the intended recipient of e-mail messages addressed to that particular address. To do so, she may contact Verisign and obtain a Digital ID for secure e-mail or establish some other form of widely accepted *public* trust relationship. For example: through PGP. She then converts "My First Soul" into digital format by generating an MP3 file.

Before Alice makes her content available through the Internet, she hashes the content, title[33], her identity and e-mail address to produce a unique hash value. She then writes a message claiming authorship of the works and signs it. This will be called the *claim*. In other words, Alice claims that the hashed content is her own by signing a message claiming authorship in which the hash of the content is present. She then obtains a timestamp on this claim. Once Alice obtains the timestamp, Alice can embed the identity certificate, the claim and the timestamp into the MP3 file[34] and make it available through the Internet. The claim may look similar to this:

---

[29] Bank checking account and routing numbers would also work with the correct implementation.

[30] such as http://www.itconsult.co.uk/stamper.htm

[31] This service is currently provided by companies such as Verisign (cost US$15/year as of Apr '03)

[32] *i.e.* a contact address

[33] We will later argue that the title should be viewed as the search criteria to be used in a peer-to-peer network to search for the file.

[34] As mentioned previously, there is a tagging mechanism (ID3 tags) for MP3 files that may be used for this purpose

"I:
Alice
with contact address:
alice@aliceworks.com
hereby claim sole authorship of the song, lyrics, musical composition and arrangement in the enclosing file and which I license for use according to:
Creative Commons license number 45
The contents of my work, its title, my identity and my contact address together hash to:
h = **0xADADFAE345344534563ED78DE5E**
under the MD5 cryptographic hash function."

(This message is signed by the Key in the digital certificate.)

The hash value is obtained by[35]:
h = MD5(Alice || alice@aliceworks.com || My First Soul || content)

In summary, the entire MP3 file would look like the following:

5.3 Obtaining Content

To obtain content from a peer-to-peer network a client usually executes a query. The result of this query may be a single file, many files or no files at all; depending on how successful the search was. If the query is successful in finding many files or many replicas of the same file, many "query hit" messages will return. We will call the set of files in all such messages the *query hit set*. In the proposed system, these "query hit" messages should include the *contact addresses* for each file matching the query. Obviously, this will generate some overhead. But it allows the client performing the query to find conflicting claims of ownership. What to do in case such conflicts are found will be addressed in detail later. For now, we will assume that any two files with similar content have claims that contain the same contact address.

Suppose Bob is one of Alice's fans. Bob wants to obtain Alice's "My First Soul" hit single. He uses his client to query the network with the search criteria set to something like: `Artist="Alice" Song="My First Soul"`. If the search is successful, Bob downloads the song just as he would do in any of the popular peer-to-peer networks currently available and listens to it as much as he likes.

[35] This could be hardened to prevent attacks where the length of the content is changed at will.

5.4 Making Donations

Suppose that Bob really liked “My First Soul” and would like to donate some money to Alice to support her work. Bob's client software can obtain Alice's e-mail address from the claim. The file also contains a certificate issued by Verisign and Alice's own signature of the claim. The client can verify Alice is claiming authorship of the song by obtaining her public key (through the certificate) and verifying Alice's own signature on the claim. Remember that Alice signed a hash that included the content. The client can recalculate the hash and check to see if it matches the original hash; this ensures the authenticity of the work. With Alice's e-mail address at hand Bob can then send her any amount of money through his Paypal account. We will call those donations MagicThanks. Note that MagicThanks do not need to be money, but they provide a return path from the user to the author. Actually, Bob might only be willing to notify Alice that he would be interested in receiving a newsletter each time a new album is released.

If the sender of the MagicThanks message does indeed want to send money, there are at least two possibilities. Money can be sent with the MagicThanks message itself as when using Paypal, or the MagicThanks may itself be only a note stating the user's intent, which may later be claimed. In the latter, the donor may be the only person able to perform some “magic” and turn the MagicThanks into real money.

5.5 Asserting Authenticity

The major functionality provided by the system is the ability to differentiate between two conflicting claims of authorship. In the present status quo, there is no certifying authority that is able to certify the authenticity of such a claim digitally. Furthermore, it is unclear whether the introduction of such an authority is desirable. The system allows a user to differentiate between two claims through the use of timestamps. All claims need to have associated timestamps to certify their timeliness. The underlying assumption is that the author herself will be the first person to have access to a completed piece of work. Thus, an author is able to provide the earliest claim of authorship and later (conflicting) claims are assumed to be illegitimate[36].

Users may want to ensure the authenticity of content for its own sake, but checking for authenticity may also be viewed as a strategy to assert quality. No one likes low quality fake content. We believe that authenticity should be viewed in this light. Because we want authenticity to be intimately tied in with the ability to detect fraud, the hash of the content generated and signed by the author should include the author’s identity[37]. This allows the user to check for authenticity, only if he is also aware of the correct identity of the author. Moreover, this also prevents a malicious party from obtaining the earliest timestamp on a work by intercepting communications between the time stamping service and the author.

A user may be unwilling to spend time and energy to inform an author that someone else is appropriating money that rightfully belongs to the author. But users may be a lot more cooperative if the time and energy spent actually go towards asserting the

[36] The issue of derivative works should be addressed further, but we will leave it for now.
[37] The intent of also including the contact address is to make sure the identity is indeed unique

authenticity of the work to be enjoyed. The idea is to unite these to procedures in one. The action performed by the end user to ensure that he or she obtains authentic material should produce as a by-product a way to detect fraud and enable authors to trace the responsible parties.

When a user decides to search the network for available content, he may be doing so with at least three different perspectives. First, the user may be unsure of what content he desires and thus may be browsing the network for any files matching a broad query. A similar but not identical situation occurs if the user knows exactly what he is looking for, but does not how to find it. For example, a user may have heard a song on the radio but is unable to find out who sings it or what the title of the song is. The user's efforts of browsing through the network are not directly related to obtaining authenticity in either of these two possibilities. In some sense, having unauthorized replicas of the work in these scenarios may actually help the user find what he wants. However, a different scenario occurs if the user knows what to look for. In this case, the user will very likely perform a very specific query and many replicas of the same content will be present in the query result set. If the user is presented with the choice of where to download the desired content from, he may use that choice to assert the authenticity of the content to be obtained. In this latter case, the user's actions should also be able to detect fraud.

We argue that a criterion that may be able to differentiate between the former two scenarios and the latter is the percentage of files with the same name in the query result set. A large set of files with identical or very similar names not only indicates that they all have identical content, but also will tend to be obtained after a very specific query. Furthermore, suppose that within files of similar names conflicting claims of ownership are found. It is in this scenario that the system may automatically engage in a procedure to determine the authenticity of the content.

The three scenarios described above are far from disjoint entities, there is a large continuum of possibilities between them. Due to the fact that we have only this blurred assessment of the environment we suggest that the actions taken by the system should also be continuous in nature. Thus, we suggest that any criterion used to determine the applicability of the mechanisms used to determine authenticity and detect fraud should be employed probabilistically. Figuratively, we make the analogy between this probabilistic line of action and the dithering of a digital image, in both cases a smoother transition between extremes better represents the reality at hand. A simple implementation of this suggestion would be to trigger the authentication mechanisms with probability that depends on the criterion used. For example, if two files of identical names and conflicting claims are found the authentication mechanism may be triggered with probability 1. If the filenames differ only by the existence of the letters 's' at the end of words an authenticity check may be triggered only half the times. Finally, if the filenames differ by an extra word the authenticity checking mechanism may be triggered only 5% of the times.

5.6 Limitations

Hopefully, the evidence produced by the system to enable prosecution of malicious parties should be considerably damaging if presented in a Court of Law.

Assuming the malicious parties were appropriating MagicThanks, digitally signed fake claims must have been produced[38]. An accused party may only use the compromise of a key as a defense strategy if the key was compromised *before* the timestamp on the fake claim was obtained. Later compromises are of no significance if the time stamping service is reliable. The original claim also binds the author to a specific licensing agreement and thus protects the user from later changes of heart. The dependability of the evidence produced however is dependent on how tightly the keys that signed the claim are tied to a real word individual or organization. If in practice the proposed solution of using e-mail certification turns out to be inadequate further measures for establishing this link may be required.

Problems may arise if an artist changes her contact address, in this situation revocation of contact information would be desirable and that requires refinements of the proposed scheme. Note that in the proposed scheme the e-mail certificates are attached to an identity and expire after a fixed period of time. To make donations with assurance after the certificate expires, it would be necessary to trace that identity. Fortunately, the same idea can be used to revoke old contact addresses if a new certificate can be found online[39].

Note that the likely profit from appropriation increases with the number of file replicas produced, but the risk of detection also increases. This can be subverted if the file replicas produced by a malicious party have completely different search criteria from the originals. Fortunately, this also implies that it will be very difficult for interested users to find the fake content generating a negative impact on the return expected by the malicious party.

The introduction of MagicThanks in a real world situation must be made with caution. Any content that is widely available before the introduction of the system runs the risk of being "hijacked" by parties able to obtain the first timestamp on a claim of authorship. Because these are publicly available, everybody – not just the author – has an equal chance of obtaining such a timestamp. Before the implementation of the system all content creators should be informed and have timestamps arranged.

6. Concluding Remarks and Future Research

The present work identifies new business models that could be used to fund musicians in the new technological paradigm presented by P2P networks. It also points out what mechanisms are available and which still need to be developed to enable such business models. A simple system to link authors of digital works to content is proposed.

Quite a few different research directions may be followed from this point. One future research direction is the implementation of MagicThanks. This is a reasonably clear path to follow and it should help to mature the ideas behind the system.

Another topic that can be addressed is the development of an open payment system capable of supporting the Distributed Patronage Model and thus, able to handle a large number of small revocable donations efficiently. Also, most electronic payment

[38] How to tie the electronic identity of the parties to their identities in the real world is a matter that should be addressed differently for each implementation.

[39] Other approaches to revocation are also possible.

systems proposed have been closed in the sense that only one entity gives financial backing to and operates the electronic currency available. For example, only Paypal Inc. runs the Paypal system[40]. This is the same paradigm used with real cash and central banks when the gold standard was in place. However, it is very different from what happens with credit cards where each bank is financially responsible for backing the cards it issues and a different agent (such as VISA or MasterCard) operates the network. A more flexible approach where more than one backing entity is in charge and may participate in the network may be preferable also for other forms of electronic payment. It would also be of value to have the payment system be able to evolve to handle future needs. These top-level requirements tend to be secondary to others in many of the proposed systems.

Other ideas or paths which are not yet so clear would be to develop the protocol to use in a *publicly verifiable* implementation of SPP or to develop an open promotion system to be used in conjunction with P2P networks. Public verifiability would also be desirable for the electronic donation billboards mentioned in the text. Moreover, these may require mechanisms to exclude some donations (such as those from sponsors).

[40] Paypal now operates with pass-through FDIC insurance.

7. References

[0] See http://www.serci.org/abajo(biblio).html

[1] K. Aberer, M. Hauswirth, "An Overview on Peer-to-Peer Information Systems", *Proceedings of Workshop on Distributed Data and Structures (WDAS-2002)*, Paris, France, 2002.

[2] M. Adler, "Stardom and Talent", *The American Economic Review*, Vol. 75, issue 1, pp 208-212, Mar. 1985.

[3] See http://www.asp-shareware.org/users/history-of-shareware.asp

[4] See http://www.altnet.com

[5] See http://www.amazon.com/honor/

[6] N. Asokan, P. A. Janson, M. Steiner and M. Waidner, "State of the Art in Electronic Payment Systems", *Computer*, vol 30, n9, pp 28-35, 1997. (showing signs of aging)

[7] N. Asokan, V. Shoup and M. Waidner, "Optimistic Fair Exchange of Digital Signatures", *IEEE Journal on Selected Areas in Communications*, vol 18 (4):593-610, 2000.

[8] M. Avalon, "Confessions of a Record Producer: How to Survive the Scams and Shams of the Music Business", 2 Ed., Backbeat Books, Feb 2002.

[9] M. Bagnoli and B. Lipman, "Private provision of public goods can be efficient", *Public Choice*, **74**:59-78, Kluwer, 1992.

[10] N. Behling , "Cyberpayments - Credit cards are here to stay", *JurPC Web-Dok. 16/2001, Abs. 1 – 32*, January 2001. http://www.jurpc.de/aufsatz/20010016.htm

[11] P. Belleflamme, "Pricing Information Goods in the Presence of Copying", Department of Economics, Queen Mary, University of London, Aug 2002. http://www.qmw.ac.uk/~ugte186/research.htm

[12] G. Bézard, "Innovation in Internet Payments: The Plot Against Credit Cards", Celent Report, April 16, 2002. Abstract available at: http://www.celent.com/PressReleases/20020415(3)/InternetPaymts.htm

[13] J. Camenisch, V. Shoup, "Practical Verifiable Encryption and Decryption of Discrete Logarithms", CRYPTO 2003, Santa Barbara, USA, Aug 17-21, 2003.

[14] See http://www.eff.org/Privacy/Digital_money/david_chaum.biblio

[15] D. Chaum, A. Fiat, & M. Naor, "Untraceable Electronic Cash," *Advances in Cryptology CRYPTO '88*, S. Goldwasser (Ed.), Springer-Verlag, pp. 319-327.

[16] See http://www.creativecommons.org/

[17] A. Duchêne, P. Waelbroeck, "Welfare Implications of Illegal Copies: the Case of Peer-to-Peer Distribution Technologies", Innovations and Intellectual Property, Brussels, Belgium, Nov 22-23, 2001. http://homepages.ulb.ac.be/~pwaelbro/

[18] P. Fishburn, A. Odlyzko and R. Siders, "Fixed fee versus unit pricing for information goods: competition, equilibria, and price wars", *First Monday*, vol. 2, no. 7 July 1997, http://www.firstmonday.dk/

[19] A. V. Goldberg, J. D. Hartline, A. Wright ,"Competitive Auctions and Digital Goods", Symposium on Discrete Algorithms, pp 735-744, 2001.
http://citeseer.nj.nec.com/goldberg99competitive.html

[20] S. Goldwasser, S. Micali, C Rackoff, "The Knowledge Complexity of Inter- active Proof Systems". *SIAM Journal of Computing*, 18(1):186-208, February 1989.

[21] C. Grothoff, "GNUnet – An Excess Based Economy", *WIRTSCHAFTSINFORMATIK*, Peer-to-Peer (P2P): technologies, architectures and applications, Issue 3/2003. http://www.gnu.org/software/GNUnet/

[22] P. Harisson, "The Rational Street Performer Protocol",
http://www.logarithmic.net/pfh?action=random

[23] A. Herzberg, "Micropayments", pp. 245-280, chapter 12 in *Payment Technologies for E-Commerce*, Ed. W. Kou, Springer-Verlag, 2003.

[24] J. L. Hougaard, K Lindeneg and M. Tvede, "Selling Digital Goods on the Internet", University of Copenhagen Oct 2002, http://www.econ.ku.dl

[25] S. Jarecki, A. M. Odlyzko, "An Efficient Micropayment System Based on Probabilistic Polling". *Financial Cryptography '97, Lecture Notes in Computer Science*. Hirschfeld, R. (ed.) Vol. 1318. Springer-Verlag, (1997) 173–191.
http://www.dtc.umn.edu/_odlyzko/doc/recent.html

[26] See http://www.kazaa.com/us/help/glossary/participation_ratio.htm

[27] J. Kelsey and B. Schneier, "The Street Performer Protocol", *The Third USENIX Workshop on Electronic Commerce Proceedings*, USENIX Press, November 1998.
http://www.counterpane.com/street_performer.html

[28] S. Liebowitz, "Policing Pirates in the Networked Age", *Policy Analysis*, 438. May 15, 2002. http://www.cato.org/pubs/pas/pa-438es.html

[29] R. J. Lipton, R. Ostrovsky "Micropayments via Efficient Coin-Flipping". *Financial Cryptography 1998*: 1-15.

[30] M. Mauss, *The Gift – Forms and Functions of Exchange in Archaic Societies*, English translation, Cohen & West, London, 1954.

[31] S. Micali , R. L. Rivest, "Micropayments Revisited", *Proceedings of the Cryptographer's Track at the RSA Conference 2002*, Bart Preneel (ed.), Springer Verlag CT-RSA 2002, LNCS 2271, pages 149--163.

[32] M. Mussa, S. Rosen, "Monopoly and product quality", *Journal of Economic Theory*, 18, pp 301-317, 1978

[33] A. Neus, "Blender and the Street Performer Protocol: Freak success or first of a trend?" http://www.blender.org/modules/bc2002/Neus-Blender-SPP.pdf, 2002.

[34] A. M. Odlyzko, "The case against micropayments," *Financial Cryptography: 7th International Conference, FC 2003*, R. N. Wright, ed., Lecture Notes in Computer Science #2742, pp. 77-83, Springer, 2003.
http://www.dtc.umn.edu/~odlyzko/doc/complete.html

[35] See www.peppercoin.com

[36] T. Regner , "Innovation of Music" , in *The Economics of Copyright: Developments in Research and Analysis*, edited by R. Watt, chap 6, pp 99-111, Cheltenham, Edward Elgar, 2002. http://www.berlecon.de/output/events/files-bite5/Regner2.pdf

[37] R. L. Rivest and A. Shamir, "PayWord and MicroMint--Two Simple Micropayment Schemes", *CryptoBytes*, vol 2, num 1 (RSA Laboratories, Spring 1996), 7--11.

[38] J. Seabrook, "The Money Note", *The New Yorker*, July 7th, 2003.

[39] A. Shamir, "How to share a secret", *Communications of the ACM*, vol. 22 , no. 11, pp. 612-613, November 1979.

[40] C. Shapiro, H. Varian, "Information Rules: A Strategic Guide to the Network Economy", Harvard Business School Press, Nov 1998.

[41] N. van Someren, "The Practical Problems of Implementing MicroMint", *Financial Cryptography 2001*: 41-50.

[42] See http://tipster.weblogs.com/

[43] H. Varian, "Buying, Sharing and Renting Information Goods", Dec 1994.
http://sims.berkeley.edu/~hal

[44] H. Varian (2001), Mail to Bruce Schneier and Chris Rasch regarding the Street

Performer Protocol and software completion bonds.
http://www.livejournal.com/talkread.bml?journal=crasch&itemid=68315#cutid1

[45] J. K. Willcox, "Where Have All the CDs Gone?", *Sound&Vision Magazine*, Aug 29, 2003. http://www.shorl.com/bemygesomove

[46] K. Yoon, "The Optimal Level of Copyright Protection," *Information Economics and Policy*, 14(3), pp. 327-348, September 2002.
http://econ.korea.ac.kr/~kiho/