Online Content Trading: Sharing Mart System and Auction Experiments

Hazer Inaltekin*, Necati E. Ozgencil †, Hao Xu†, Mung Chiang†, H. Vincent Poor†
* Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, Australia
† Department of Electrical Engineering, Princeton University, Princeton, NJ, USA

***Abstract*—This paper introduces an operational monetary system, the Sharing Mart system, for online content trading, and discusses what it can do with a toy experiment. Sharing Mart is a *virtual-money-based* file sharing system in which different digital rights (e.g., *view only*, *download* and *resell* rights) of various file types (e.g., *video*, *audio*, *graphics* and *documents*) can be traded by means of different transaction styles (e.g., *marked price transactions* and *S-Mart auctions*). Sharing Mart is the content equivalent of stock exchange markets, eBay and Amazon. The fully operational version has been used by 250 students in the School of Engineering and Applied Science at Princeton University. After a quick introduction to Sharing Mart, its capabilities as a testbed are illustrated by means of a multiple-winner content auction experiment. Empirical data indicate that revenue earned by selling content and probability of winning an auction are maximized as seller/buyer strategies approach the Nash equilibrium. Content sellers and buyers lacking of strategic thinking lead to revenue reductions and reductions in the success rate to win an auction, respectively.**

I. Introduction

The world has witnessed tremendous growth in both communication and social networks over the last decade. Nowadays many people around the globe, from teenagers to elderly people and terrorists, are ubiquitously connected both physically and virtually. The viral growth of social networks is particularly impressive: Facebook currently has more than 150 million active users and has an average of $250,000$ new registrants per day, with an average of 3% weekly growth since January 2007. This puts the networks, either social in virtual online space or technology networks among communication devices, at the center of our society. Therefore, it is of particular importance to understand and reveal under-explored functional and topological interactions between social and technological networks.

To this end, fundamental research is promising to reveal the structure of complex interconnected networks, and various distributed processes running on them [1], [2] and [3]. Mathematical models developed over the last decade have also been justified with real human data to some extent, however, there still remains an essential need to collect large amounts of real human data to develop these models further, to tap into complex interactions between social-overlay/technology-underlay networks, and to advance the current knowledge base in the field of socio-technological networks research. Such real human data is available on the Web today, and it can be collected from various online social networks [4], [5] and [6]. However, the biggest bottleneck with this approach is the excessive amount of noise present in this data, which makes the distillation of the data describing and explaining desired interactions from the noisy data a prohibitively difficult task.

An alternative is to make use of advances in Web technologies and collect such data in a controlled manner by designing virtual online experiments involving many human subjects [7], [8] and [9]. In such experiments, we have the flexibility of controlling several design degrees of freedom, and understanding their individual or collective effects on the final emerging global network behavior by adjusting them one at a time or in groups. This further enables researchers to analyze rare events, large groups and interacting factors in socio-technological networks by means of online social experiments.

This work introduces an *online content trading platform*, which we call Sharing Mart (S-Mart), developed recently at Princeton University to perform controlled experiments in social file sharing systems constructed on top of technological networks as overlays. S-Mart is a *virtual money* based file sharing system, in which different digital rights (e.g., *view only*, *download* and *resell* rights) of various file types (e.g., *video*, *audio*, *graphics* and *documents*) can be traded by means of different transaction styles (e.g., *marked price transactions* and *S-Mart auctions*). S-Mart is envisioned as being a content equivalent of stock exchange markets, eBay and Amazon. Further details of the proposed system are provided in Section II. The current system [10] has been used by 250 students in the School of Engineering and Applied Science at Princeton University. We are also willing to open part of the developed operational system to researchers interested in experimenting in file sharing systems with human subjects, similarly to what has been done with Planet Lab for communication network researchers.

S-Mart's monetary incentives together with user ratings eliminate the free rider problem in file sharing systems. In S-Mart, end-users determine prices of their content items in terms of a virtual currency called *tokens* based on the time and effort needed to produce them. Therefore, initial file prices in S-Mart reflect subjective evaluations of their producers. We then let open market dynamics, competition from similar content, and supply and demand forces determine final content prices.

S-Mart is more than simply a new approach to file sharing. It also serves as a powerful testbed to perform controlled experiments with human subjects to better understand human behavior in exchanging content under the guidance of pricing signals, and other sociologically complex phenomena such as time evolution of the popularity of a content item or a user. To illustrate its usage as a testbed, we designed a toy auction based competition among seven graduate students in the Electrical Engineering Department at Princeton University. The purpose of this auction experiment is to understand strategic behavior (or the lack thereof) of users in online auctions. We solicited students to submit buyer and seller strategies to trade files through S-Mart. Then, the submitted strategies competed against each other to garner the maximum number of points, which points were a combination of total seller revenue and the number of auctions won. We observed that the revenue earned by selling content and the probability of winning an auction increase as seller and buyer strategies approach the socially optimal Nash equilibrium of the designed experiment. Further details of the experiment and initial results are provided in Section III.

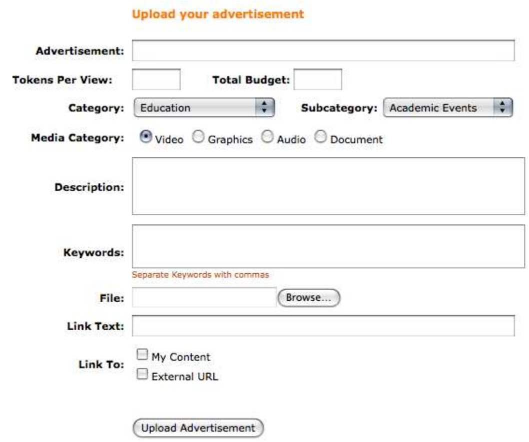

Fig. 1. User interface for uploading ads.

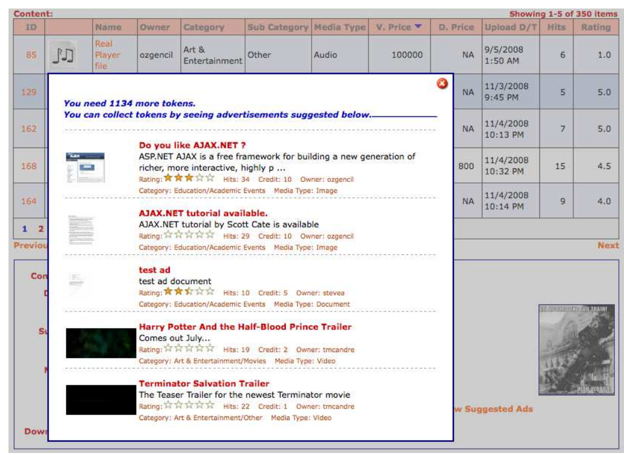

Fig. 2. S-Mart's targeted ad mechanism.

II. Details of the Sharing Mart System

In this section, we explain the details of the selected properties of the proposed virtual money based file sharing system. We start with a description of the ramping up of the initial user activity in S-Mart.

A. Ramping up Initial User Activity

The initial ramp up of user activity in S-Mart was achieved by providing each new S-Mart member with initial start-up funding (in terms of tokens) with which to begin file trading. Therefore, the total amount of virtual money in the system increases with user involvement, and more content trading activities are expected to appear with an increasing number of tokens in the system. However, some users' token balances may decrease over time if they cannot attract enough interest in their content items from others. We support such users by means of S-Mart's advertising mechanism, whose details are explained in the next section. Furthermore, such users can also buy resell rights of popular content items in S-Mart, and start contributing to S-Mart's ecosystem by sharing these files with others. This creates a liquid secondary market for content similar to stock exchange markets, with increased availability of the same content from different users. In this ecosystem, some content brokers are expected to appear that buy resell rights to content when it is unpopular, and sell them for higher prices when it becomes popular. We also reward users with tokens for their positive contributions to S-Mart such as rating and reviewing files that they have bought.

B. S-Mart Advertising Mechanism

In the S-Mart ecosystem, there are two kinds of content items: **(a)** positive-valued content items and **(b)** negative-valued content items. Positive-valued content items are regular content items for sale whose values are paid to content producers by content consumers. Negative-valued content items are advertisements whose values are paid to content consumers by content producers. An S-Mart user can upload both types of content items to S-Mart. To upload an ad, a brief description of the ad, an ad budget, number of tokens to be paid per ad view, and the ad type need to be entered. Ads can also be targeted to a specific group of users by using S-Mart's group formation mechanism. The user interface for uploading ads is shown in Fig. 1.

The S-Mart advertising mechanism has rewards for both content producers and content consumers. By means of this ad mechanism, content producers are able to advertise their content to other members. Content consumers can consume the ads recommended based on their empirical user profiles whenever they are short of tokens to buy their desired content items. Figure 2 illustrates the targeted ad action. Examples of empirical human data for user profiling include users' past interests (e.g., content type and category), users' past bidding and buying patterns (e.g., the percentage of their net wealth used while buying a content item) and social graph data (e.g., social connections of users and what their friends are doing). Thus, this monetary system for online content trading sharpens user profiling.

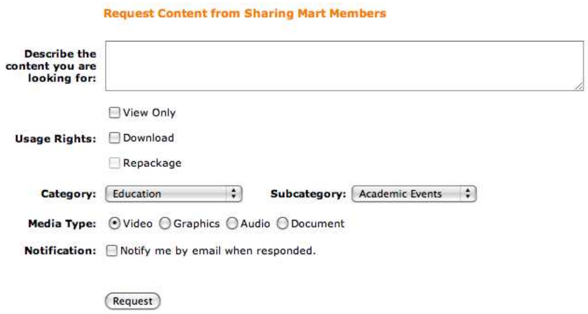

Fig. 3. User interface for requesting files.

C. *What Can Users Do in Sharing Mart?*

There are two main types of users in S-Mart: **(a)** content consumers and **(b)** content producers. Content producers are the catalytic users with creative minds who are able to create interesting content items that attract attention from other users. Content producers can also be content consumers when they want to buy others' files. These users usually do not have any difficulty in buying others' files since they typically have enough tokens in their bank accounts.

Content consumers are the users who cannot contribute interesting popular content to the system. We support such users by means of start-up funds, the ad mechanism, the token rewarding mechanism and the secondary content market mechanism, as mentioned above.

Content consumers can spend their tokens to buy various digital rights to content items at different prices via different transaction styles. S-Mart differentiates among view-only, download and resell rights of content items. These rights can be traded by using either marked price transactions or multiple-winner S-Mart auctions. If content consumers cannot find files that they are looking for, they can request them by using S-Mart's content request mechanism. To request a file, a content consumer enters its brief description, file type (document, audio, video ant etc.) and the tokens offered to access its different digital rights. If a content producer responds to the content request, the requester is notified via e-mail, and she can access these responses whenever she logs into the system. The user interface for requesting files is shown in Fig. 3.

S-Mart allows content producers to design their content stores, and trade different digital rights at different prices by means of different transaction styles. We will mention the details of S-Mart's auction mechanism in Section III. Content producers can also earn tokens by responding to content requests.

D. *Engineering Design and Incentive Mechanism*

S-Mart consists of two main layers: **(a)** a data plane and **(b)** a control plane. The data plane can be implemented by using either peer-to-peer (P2P) technology or a classical server-client technology. The control plane is implemented by means of a central secure server. The data plane is responsible for data transfer and storage. The control plane is responsible for lightweight system control related functions. The current implementation uses client-server technology at the data plane. Such a two-layered system design is more scalable in terms of potential future innovations.

The control plane serves four main purposes: **(a)** as a central bank (i.e., it adjusts the total token amount in the system), **(b)** as a commercial bank (i.e., it updates user accounts after each file transaction and generates monthly trading reports), **(c)** as a police station (i.e., it preserves different digital rights via watermarking) and **(d)** as an investment bank (i.e., it enables a secondary market).

S-Mart's incentive mechanism consists of monetary incentives and user ratings. Monetary incentives eliminate free-riders, and the currency-based system provides anyone - anytime file trading capabilities for users. User ratings deter users from uploading corrupted or junk files.

E. *Some Other Capabilities*

Some other functions of S-Mart are group formation, internal messaging, blogging and detailed reports about user statistics. Further details about them can be found in the S-Mart user manual [11].

III. S-MART AS TESTBED: AN AUCTION EXPERIMENT

The classical method used to discover the price of an item when a seller is uncertain about it is to let others submit their bids for this item (see [12], [13], [14]). With this motivation, we have built an auction platform to enable content producers to solicit bids from content consumers to discover fair market prices for their content items. In this part of the paper, we first describe the S-Mart auction platform and the multiple-winner S-Mart auction. Then, we report our initial results of a toy auction experiment with human subjects for online content trading.

A. *S-Mart Auctions*

Selling a file is different from selling a tangible good via an auction mechanism because multiple copies of a file can be sold in a single auction. Therefore, we implemented the following *uniform price*, *unit demand* and *multiple-winner* file auction in S-Mart. The seller enters the minimum price, auction start/end date and time, and digital rights and the number of copies of the file, M, to be sold. Then, the file is sold to the highest M bidders at the price of the $(M+1)^{\text{st}}$ highest bid. If the number of unique bidders is smaller than M, the file sale price becomes equal to the minimum price, and all bidders pay this minimum price. We first prove an important property of this auction.

Definition 1: An auction mechanism is said to be **incentive compatible** if it induces each bidder to submit a bid that sincerely reflects her true value for the item.

An incentive compatible auction is also an efficient (i.e., items are allocated to those who most value it) and standard (i.e., items are allocated to the highest bidders) auction.

Theorem 1: S-Mart auctions are incentive compatible.

Proof: Let v_i be i^{th} bidder's true valuation for the file, b_i be her bid, and $b^{(M+1)}$ be the $(M+1)^{\text{st}}$ highest bid at the end of the auction. First consider $v_i \geq b^{(M+1)}$. Submitting $b_i > v_i$ does not increase the bidder's utility (i.e., $v_i - b^{(M+1)}$) when compared with submitting a bid $b_i = v_i$. On the other hand, she can lose the item with positive probability if she submits a bid $b_i < v_i$. As a result, bidding $b_i = v_i$ is a weakly dominant strategy for this case. The same is true for $v_i < b^{(M+1)}$. ■

An alternative for S-Mart auctions is discriminatory multi-unit auctions (i.e., the highest M bidders get the file by paying their bids) However, the revenue equivalence principle (see [13]) states that the seller revenue tends to stay the same for different auctions when some mild conditions are met. Therefore, we have decided to implement S-Mart auctions in the current operational version of the S-Mart system.

Content producers selling their files via S-Mart auctions now face a complex stochastic revenue maximization problem over the optimization variables such as the number of copies of a file to be sold, minimum asking price and the auction duration. This optimization problem will be subject to our auction experiment in the next section.

B. Auction Experiment Set-up

Our experiment was a 96-hour agent-based competition among seven graduate students at Princeton as a part of a homework problem set. Our subjects were given a lecture on game theory and auction theory before the experiment. Therefore, they were familiar with strategic thinking in a game setting.

We formed a closed test group among seven students by using S-Mart's group formation function, and gave each student 1800 tokens for file trading. Each student was requested to propose a seller strategy, and to set three S-Mart auctions according to his or her selling strategy. Each student was also requested to submit an automated bidding agent to bid for 18 other auctions. The start time/date of all auctions was the same, but the end time/date varied according to seller strategies with the restriction that all auction durations were less than 96 hours.

Student i, $1 \leq i \leq 7$, got 100 points from each auction s/he wins, and the total number of points p_i that student i got was equal to the total revenue obtained from three auctions (in terms of tokens) plus 100 times the number of auctions won. Therefore, p_i is equal to

$$p_i = \sum_{k=1}^{3} f_i^k M_i^k + 100 \cdot \sum_{j \neq i}^{7} \sum_{k=1}^{3} \delta_{j,i}^k, \tag{1}$$

where f_i^k is the final price of the k^{th} auction from student i, M_i^k is the number of copies in the k^{th} auction from student i, and $\delta_{j,i}^k$ is an indicator function that equals 1 if student i wins the k^{th} auction from student j, and equals 0 otherwise. Note that p_i is a complex function of M_i^k, the minimum sale price R_i^k, the auction duration T_i^k, the bidding strategy $\beta_i(t)$, $t \in [0, 96]$, of student i and the bidding strategies $\beta_j(t)$ of the other students. In addition to the time t, the bidding strategies also depended on other factors such the minimum and current price of the file, which we do not show for notational simplicity.

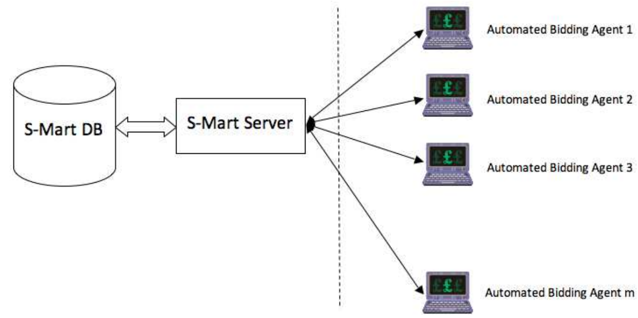

Fig. 4. Automated bidding agents.

The final score of student i is calculated by normalizing his or her point total with the maximum number of points. As a result, student i's objective function to maximize is

$$s_i = \max_{\substack{\{M_i^k\}_{k=1}^3, \{R_i^k\}_{k=1}^3, \\ \{T_i^k\}_{k=1}^3, \beta_i}} \frac{p_i}{\max\{p_j : j = 1, 2, 3, \cdots, 7\}} \tag{2}$$

given the bidding and selling strategies of other students.

Note that leftover tokens at the end of the competition are not counted towards p_i in (1) to prevent students from having hard upper limits (i.e., 100 tokens) in their bidding strategies.

Even though the optimization problem faced by each student is complex, the experiment was set-up such that it has a simple, symmetric and socially optimal Nash equilibrium.

Theorem 2: The above experiment has a symmetric and socially optimal Nash equilibrium at which $M_i^k = 6$, $R_i^k = 100$, and $\beta_i(t) = 100$ for $1 \leq i \leq 7$, $1 \leq k \leq 3$ and $t \in [0, 96]$.

Proof: Given the seller strategies in the theorem, file prices stay constant at 100 tokens since there are 6 students competing for each auction. Therefore, a student can win all auctions by bidding 100 tokens for each of them, which is the best s/he can do as a buyer. Similarly, given the bidding strategies of the others, a student can sell 18 copies of three different files at three different auctions by setting the number of copies per auction to 6 and the minimum sale price to 100, which is the best a student can do as a seller. ■

Observe that all students collect 3600 points, and therefore a score of 100, at this Nash equilibrium.

C. Automated Auction Bidding Agents

In order to eliminate random factors (e.g., forgetting to bid for an auction) in bidding, we developed a Java library to enable automated bidding for auctions. In addition, it is hard for a student to remember all parameters of 18 ongoing auctions and to bid on time manually. Java classes developed and their functions are shown in Table I. Students use these classes to program their automated bidding agents. These bidding agents run on students' computers, connect to the S-Mart server, retrieve auction information to decide on a bid amount, and then submit the bids as shown in Fig. 4.

D. Empirical Results

In this section of the paper, we report empirical findings about the auction experiment. We start our analysis with

TABLE I
AUTOMATED BIDDING AGENT LIBRARY

Class Name	Function
AuctionInfo.java	Retrieves auction information such as number of copies, minimum price, current price, start/end date/time, content ID and seller ID from the S-Mart database.
Bidder.java	Submits bids to the server.
BalanceInfo.java	Retrieves the bidder's current balance information from the S-Mart database.
UserProperties.java	Sets user properties for secure authentication.

seller strategies, and observe that seller revenues of students are maximized as their seller strategies approach the Nash equilibrium seller strategy in Theorem 2.

The parameters for revenue optimization as a seller in S-Mart are the number of copies of the file, the minimum sale price, and the auction duration. Students employed different seller strategies by varying each one of these parameters. A brief summary of student strategies is as follows.

- *Auction Durations:* Three students spread their auction end times over 96-hour period to gather funds earlier, and then to use them for bidding in other auctions. The others set their auction durations to 96 hours.
- *Minimum Prices:* Six students realized that the reward for winning an auction is 100 points, and therefore set their minimum prices to be less than 100 to attract enough interest from others. This created active bidding wars among bidders. One of them set different minimum prices, 70, 100 and 150, to explore weaknesses and the lack of strategic thinking in bidder strategies.
- *Number of Copies:* Four students tried to exploit the trade-off between the number of copies sold and the final sale price by setting the number of copies to small numbers in the range of one to four with the hope that final price would go sufficiently up that the effect of selling a small number of copies would be compensated by high final sale prices. Three students employed a strategy close to the Nash equilibrium seller strategy proved in Theorem 2 by setting the number of copies to 5 and 6.

We analyze the effect of each one of these parameters on the final seller revenue by starting with the most interesting parameter: number of copies of a file to sell.

When averaged over the best three seller strategies, we found that the average number of copies maximizing the seller revenue is 5.33, which is close to the 6 copies result suggested by Theorem 2. This further motivates us to ask the following question: what happens if all students set the number of copies to be sold to 6 by keeping other parameters (e.g., auction duration and minimum price) constant? In this case, the file sale price stays the same at the minimum price set in the beginning of an auction, and the total revenue equals the number of unique bidders times the minimum price. Our findings are summarized in Fig. 5. Except for one student, revenue of a seller summed over three auctions strictly increased. The average revenue increase is 20%. The student whose revenue stays the same already set the number of copies

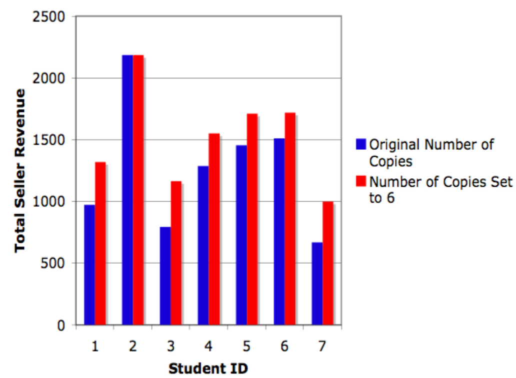

Fig. 5. Change of seller revenues when the number of copies to be sold is set to 6.

per auction to 6. These findings are another justification for employing the Nash equilibrium seller strategy to maximize seller revenue.

We can further analyze the average revenue per seller as a function of the number of copies by using the statistics in Table II regarding the relationship between the number of copies, the final sale price and the percentage of copies sold. As shown in Fig. 6, the seller revenue peaks when the number of copies is around 5 to 6, which is also in accordance with the Nash equilibrium seller strategy. We expect that student seller strategies would approach the Nash equilibrium seller strategy more in terms of setting both the minimum price and the number of copies per auction with one piece of extra information: tell them before the experiment that empirical data shows that setting the number copies per auction around 6 maximizes seller revenue. This point is subject to further study but clearly shows the power of information and learning in game theory and auction theory to design systems at equilibrium.

We next quantify the effects of the minimum price on the final seller revenue. We first observe that the average minimum price over all 21 auctions was 96.67 tokens, and the median minimum price over the three best seller strategies was 95 tokens, which are close to the numbers suggested by the Nash equilibrium in Theorem 2. Secondly, we observe that a high minimum price decreases the percentage of copies sold (i.e., all copies are sold in all auctions if the minimum price is less than 80 tokens but none of the auctions have all copies sold if the minimum price is greater than 100 tokens); a higher minimum price increases the final price in all auctions, and also increases the total seller revenue in 85.71% of all

TABLE II
RELATION BETWEEN THE NUMBER OF COPIES, THE FINAL SALE PRICE AND THE PERCENTAGE OF COPIES SOLD

Number of Copies	Final Price / Minimum Price	Percentage of Copies Sold
1	231.58%	100.00%
2	205.00%	100.00%
3	118.27%	100.00%
4	103.14%	95.00%
5	102.09%	97.14%
6	100%	71.43%

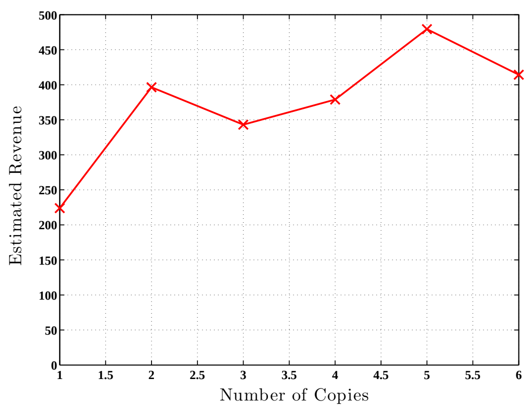

Fig. 6. Change of seller revenues as a function of the number of copies. The minimum price is set to 96.67 tokens, which is the average of minimum prices over all auctions.

auctions. Empirical data also shows that a higher auction duration increases the final file sale price. The average increase in the final price is 20.85% of the minimum price if the auction duration is between $64 - 96$ hours. This decreases to 1.58% and 1.39% for auctions whose durations are between $32 - 64$ hours and $0 - 32$ hours, respectively.

We briefly mention our empirical findings regarding buyer strategies. As already observed in the literature before (see [12]), many experienced bidders in eBay use the sniping strategy (i.e., bid at the very last minutes). The three students with the highest success rates snipe within the last 60 seconds of auctions. Their average success rate is 96.87%. That is, they win 96.87% of all auctions in which they participated. Furthermore, average number of bids per auction from the three most successful bidder strategies is 1.316, which is in compliance with the incentive compatibility property in Theorem 1: *just bid once with your true value*. Other students employ continues bidding strategy, and withdraw when the current price reaches to their true value of the file. Their success rate reduces to 61.87%, and they bid 8.625 times per auction.

IV. CONCLUSIONS

This paper serves two purposes. The first one is an introduction to a fully operational monetary system, the Sharing Mart system, for online content trading. It has been used by 250 students in the School of Engineering and Applied Science at Princeton University. In Sharing Mart, different digital rights (e.g., *view only*, *download* and *resell* rights) of various file types (e.g., *video*, *audio*, *graphics* and *documents*) can be traded by means of different transaction styles (e.g., *marked price transactions* and *S-Mart auctions*). The second purpose is to present opening discussions regarding the capabilities of this platform as an experimental testbed in file sharing systems. To this end, a toy auction experiment and some initial empirical findings backed-up with game theory and auction theory have been presented. Selected parts of the developed operational system can be opened to researchers interested in experimenting in file sharing systems with human subjects.

V. ACKNOWLEDGEMENTS

This work was supported in part by the U.S. National Science Foundation under Grant CNS-09-05086 and in part by the Princeton EDGE Lab.

REFERENCES

[1] D. Watts and S. H. Strogatz, "Collective dynamics of small-world networks," *Nature*, vol. 393, no. 6624, pp. 440-442, 1998.
[2] J. M. Kleinberg, "The small world phenomenon: an algorithmic perspective," *Proceedings of the 32nd Annual ACM Symposium on the Theory of Computing*, New York, 2000.
[3] H. Inaltekin, M. Chiang and H. V. Poor, "Average message delivery time for small-world networks in the continuum limit," *Proceedings of the 2008 IEEE International Symposium on Information Theory*, Toronto, Canada, July 2008.
[4] D. L.- Nowell and J. Kleinberg "Tracing Information flow on a global scale using Internet chain-letter data," *Proc. National Academy of Sciences*, vol. 105, no. 12, pp. 46334638, March 2008.
[5] K. Lerman and A. Galstyan, "Analysis of social voting patterns on Digg," *Proceedings of the First ACM SIGCOMM Workshop on Online Social Networks*, Seattle, WA, August 2008.
[6] M. Cha, A. Mislove, B. Adams and K. P. Gummadi, "Characterizing social cascades in Flickr," *Proceedings of the First ACM SIGCOMM Workshop on Online Social Networks*, Seattle, WA, August 2008.
[7] P. S. Dodds, R. Muhamad and D. J. Watts, "An experimental study of search in global social networks," *Science*, vol. 301, no. 5634, pp. 827-829, August 2003.
[8] M. J. Salganik, P. S. Dodds, and D. J. Watts, "Experimental study of inequality and unpredictability in an artificial cultural market," *Science*, vol. 311, no. 5762, pp. 854-856, 2006.
[9] W. S. Bainbridge, "The scientific research potential of virtual worlds," *Science*, vol. 317, no. 5837, pp. 472-476, July 2007.
[10] N. E. Ozgencil, H. Inaltekin, M. Chiang and H. V. Poor, "A monetary system for trading and pricing content," *http://sharingmart.princeton.edu*.
[11] H. Inaltekin, N. E. Ozgencil, M. Chiang and H. V. Poor, "S-Mart user manual," *http://sharingmart.princeton.edu/UserManual.pdf*.
[12] K. Steiglitz, *Snipers, Shills & Sharks*, Princeton University Press, Princeton, NJ, 2007.
[13] V. Krishna, *Auction Theory*, Academic Press, San Diego, 2002.
[14] J. Lusk and J. F. Shogren, *Experimental Auctions: Methods and Applications in Economic and Marketing Research*, Cambridge University Press, Cambridge, UK, 2007.