

Hidden Markets: UI Design for a P2P Backup Application

Sven Seuken*
 School of Engineering & Applied Sciences
 Harvard University, Cambridge, MA
 seuken@eecs.harvard.edu

Kamal Jain, Desney Tan, Mary Czerwinski
 Microsoft Research
 One Microsoft Way, Redmond, WA
 {kamalj,desney,marycz}@microsoft.com

ABSTRACT

The Internet has allowed market-based systems to become increasingly pervasive. In this paper we explore the role of user interface (UI) design for these markets. Different UIs induce different mental models which in turn determine how users understand and interact with a market. Thus, the intersection of UI design and economics is a novel and important research area. We make three contributions at this intersection. First, we present a novel design paradigm which we call *hidden markets*. The primary goal of hidden markets is to hide as much of the market complexities as possible. Second, we explore this new design paradigm using one particular example: a P2P backup application. We explain the market underlying this system and provide a detailed description of the new UI we developed. Third, we present results from a formative usability study. Our findings indicate that a number of users could benefit from a market-based P2P backup system. Most users intuitively understood the give & take principle as well as the bundle constraints of the market. However, the pricing aspect was difficult to discover/understand for many users and thus needs further investigation. Overall, the results are encouraging and show promise for the hidden market paradigm.

Author Keywords

Economics, UI Design, Market Design, P2P Backup

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces; J.4 Social and Behavioral Sciences: Economics

General Terms

Design, Economics, Experimentation, Human Factors

INTRODUCTION: MARKETS AND UI DESIGN

Many people think of Amazon or eBay when they hear about *electronic markets*. They know they can buy goods on Amazon for a fixed price, or on eBay via an auction. In these systems, monetary transactions are natural and the markets are conceptually simple such that even non-expert users can

effectively interact with them. However, in recent years, we have seen the emergence of more and more non-traditional electronic markets. For example, some toll roads adjust their prices dynamically as traffic changes¹. Digital content like music files is sometimes priced variably based on demand². Recent progress on micropayment systems might soon pave the way for even more electronic markets by significantly reducing transaction costs³. While these new, non-traditional markets often provide large benefits to the users, they can also be unnatural or complex such that individuals may not have an easy time or may not enjoy interacting with them.




Figure 1. The hidden market UI wraps around the complex underlying market and exposes a simpler interface, invoking a particular mental model in the user, whose actions still directly influence the market.

The Hidden Markets Paradigm

For these new market-based systems to become ubiquitous, we must provide a method for unsophisticated users not only to interact with them, but also to build appropriate mental models. This is particularly important for environments where monetary transactions are unusual. In this paper we propose *hidden markets*, a new design paradigm that attempts to mask as much of the prices, account balances, trading constraints, etc. from the user as possible (illustrated in Figure 1). To do this, we project a “hidden market UI” wrapped around the actual market to expose a simplified interface to the user. The goal in designing this hidden market UI is to make the interaction for the user as easy as possible while still maintaining a true feedback loop between the user and the market. This feedback is necessary for optimal outcomes for the individual user and the overall market. The key to success of these markets, then, lies in our ability to design appropriate UIs using the hidden market paradigm.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10 – 15, 2010, Atlanta, Georgia, USA
 Copyright 2010 ACM 978-1-60558-929-9/10/04...\$10.00.

*This work was done while the author was an intern at Microsoft Research.

¹www.wsdot.wa.gov/projects/SR167/HOTLanes

²online.wsj.com/article/SB123126062001057765.html

³www.niemanlab.org/pdfs/Google.pdf

P2P Backup: A Market Design Problem

In this paper we instantiate our work within a peer-to-peer (P2P) backup system and its underlying market. In a P2P backup system, users who want to back up their data must in return supply four resources to the system: storage space, upload bandwidth, download bandwidth, and online time. P2P backup is an attractive alternative to server-based systems because the immense costs of large data centers can be saved by using idle resources on millions of private computers instead. Our system uses a hybrid P2P architecture where all backup data is transferred directly between peers, but a dedicated server coordinates all operations and maintains meta-data about the location and health of the files.

The P2P backup system is novel in that it uses a market to allocate resources, more efficiently than a non-market-based system could. Each resource has a price, and the relative prices reflect the relative scarcity of the resources in the system. P2P backup is a particularly good application of hidden markets for three reasons. First, this is an environment where users are usually not thinking about monetary transactions and where they do not want to think about “selling” or “buying” resources when they make backups. Thus, exposing the users to concepts like account balances, income, expenditure, prices, etc., would be unnatural and thus undesirable. Second, the underlying market is complex, with multiple resources and many constraints. Third, the target users of this backup system are not technically sophisticated and thus the UI must be simple and facilitate operations. All of this motivates the need for “hiding the complexity of the underlying market.” Thus, by studying this example in detail we learn something about designing UIs for P2P backup systems in particular, and about hidden markets in general.

The Challenge: Hidden Markets for P2P Backup

The first challenge in designing the hidden market for a decentralized backup system is the combinatorial nature of the problem. All users must provide a certain amount of all resources, even if they currently only consume a subset of them. For example, a user who only contributes storage space is useless to the system because no files could ever be sent or received from that peer if no bandwidth is provided. We call these combinatorial requirements of the market the *bundle constraints* because only bundles of resources are worth anything. Displaying the bundle constraints in a simple way is a major challenge for the UI design.

A second challenge is that all users have different preferences regarding how much of each resource they want to supply. Some users might need their own disk space a lot and thus prefer to sacrifice their internet connection. Other users might use their bandwidth for services like VOIP or file-sharing and might have a high disutility if the quality of those services were affected. Thus, a rigid accounting system that enforces the same resource ratios across all users is undesirable. We allow different users to provide different ratios of their resources, and we update prices regularly taking into account aggregate supply and demand of all resources. However, exposing these prices to the users in an intuitive way is a very challenging task, in particular because we do not want to make the users think of a monetary market.

Previous Work & Overview of new Results

For a P2P backup system to become reality, many research challenges had to be addressed. In [14] we have described the economic market design in detail and presented a theoretical analysis of the resulting economy. This paper complements that work: we purposefully leave out the formal economic market model and instead focus on the user interface to interact with the underlying market. In a workshop contribution [13] we have previously introduced the idea of designing a hidden market UI for a P2P backup system and presented an earlier iteration of the UI.

In this paper, we make three new contributions. First, we introduce a principled way of thinking about the “hidden markets” paradigm for UI design. Second, we propose a novel UI for the market underlying a P2P backup system, a real-world instantiation of hidden markets. The UI allows the users to interact with the market without specifying bid/ask prices. Instead, the users only have to move three sliders to specify how much of each resource they want to supply. A bar chart showing how much more a user can back up given the current supply summarizes that user’s account. Our third contribution is a formative usability study which was conducted to learn about users’ views of P2P backup in general and to evaluate the hidden market paradigm. Our findings from observations and questionnaires indicate that for more than half of our participants a P2P backup system would be an attractive alternative to server-based systems. Furthermore, for almost 75% of the users the higher efficiency of a market-based system has real utility in practice: they are willing to deal with a more complex UI in return for having more freedom in deciding on their own supply settings. Finally, the study showed that most users have an intuitive understanding of the market interface and can learn, over time, to solve complex tasks with the hidden market UI. This provides the first encouraging results that with careful attention to UI design, the hidden market paradigm might be viable. However, we also present a series of shortcomings and problems we found during the study and discuss resulting implications for future market and UI design.

RELATED WORK

In addition to a deep understanding of how existing markets work, recently researchers are also gaining a better understanding of how to design new ones [11]. This has given rise to a new field called *market design*. A fundamental assumption most designers of electronic markets make is that participants are sophisticated users able to specify bids in an auction-like framework (e.g., [1, 3, 7]). A particularly sophisticated UI design for a complex combinatorial market is described in [12]. However, unlike in energy markets, grid networks, or sourcing auctions, we are targeting millions of non-expert users and thus such interfaces are not practical for a P2P backup system. The application most similar to our system is *Wuala* (www.wuala.com), an implementation of a P2P backup service. However, *Wuala* does not use a market and thus we cannot learn from its UI design.

Recently, HCI researchers have gotten more interested in topics at the intersection of UI design and economics. Hsieh et al. [6] test whether the use of markets in synchronous

communication systems can improve overall welfare. Hsieh et al. [5] explore a similar idea in the domain of Q&A applications where users could attach payments to their questions. While the use of the markets is similar in vein to our approach, i.e., using markets to most efficiently allocate resources as is standard in economics [4], in both papers they used a very explicit UI showing monetary prices to the users. This is something the hidden market approach tries to avoid.

MySong [9, 15] is an application that allows musical novices to compose new songs. The authors have successfully designed and tested a UI that lets users interact with complicated elements of the underlying machine learning system in an intuitive way. Although this work does not concern the design of market interfaces, the approach is similar to ours, in the sense that both user interfaces are designed to hide the complexity of the underlying system while maintaining the important feedback loop.

Satu and Parikh [10] compare live outcry market interfaces in scenarios such as trading pits and electronic interfaces. They draw a distinction between trying to blindly replicate the real world in the UI, and locating “defining characteristics” that must be supported. In our work, we adopt this philosophy and attempt to mask the unnecessary affordances in the hopes that the relevant ones become easier to use.

The work that is closest to our approach is *Yoopick*,⁴ a combinatorial sports prediction market [2]. This application provides a very intuitive UI for trading on a combinatorial prediction market. The designers successfully hide the complexity of making bets on combinatorial outcomes by letting users specify point spreads via two sliders. This approach is very much in line with the hidden market paradigm.

A P2P RESOURCE EXCHANGE MARKET

In this section we define the key concepts regarding the P2P market design, to the degree necessary to understand the particular UI we propose. Note that the whole system as described in this paper is implemented and an internal alpha has been successfully tested.

Suppliers & Consumers

Each user in the system is simultaneously a supplier and a consumer of resources. A single peer on the consumer side demanding a service (backup, storage, or retrieval) needs multiple peers on the supplier side offering their resources (space, upload and download bandwidth, and online time). The production process of the server (bundling multiple peers and coordinating them) is essential, turning unreliable storage from individual peers into reliable storage. Note that each peer on the supplier side offers a different bundle of resources while each peer on the consumer side gets the same product, i.e., a backup service with the same, high reliability.

Minimizing Replication: Erasure Coding

One natural concern about P2P backup is that individual P2P users have a much lower availability than dedicated backup

servers. Thus, a P2P system must maintain a higher file redundancy to guarantee the same file availability as server-based systems. Simply storing multiple file copies would be very costly. Fortunately, we can significantly reduce the replication factor by using *erasure coding* (see [8] for an application of erasure coding to P2P storage). The erasure code splits up a file into k fragments, and produces $n > k$ new fragments, ensuring that *any* k of the n fragments are enough to reconstruct the file. Using this technique, we can achieve the same high reliability as server-based systems (e.g., 99.999% availability) while keeping replication low. For example, if the users are online 12h per day on average, we can achieve a replication factor as low as 3.5.

Operations in the Backup System

We consider the following five high-level operations:

1. **Backup:** When a user performs a backup, file fragments are sent from the consumer to the suppliers.
2. **Storage:** The suppliers must persistently store the fragments they receive (until they are asked to erase them).
3. **Retrieval:** When a user retrieves a backup, file fragments are sent from the suppliers to the consumer.
4. **Repair:** When the server determines a backed up file to be unhealthy, the backup is repaired.
5. **Testing:** If necessary, the server initiates test operations to gather new data about a peer’s availability.

Each of the operations requires a particular set of resources:

Operation	Resources Required from Suppliers
1. Backup	Download Bandwidth
2. Storage	Space
3. Retrieval	Upload Bandwidth
4. Repair	Download and Upload Bandwidth
5. Testing	Download and Upload Bandwidth

Table 1. Operations and their Required Resources.

Prices, Trading & Work Allocation

While our long-term goal is the design of an open P2P market using real money, here we consider a system where monetary transfers are prohibited, i.e., all trades in the market are done using “virtual currency.” Each resource has a price at which it can be traded and in each transaction the suppliers are paid for their resources and the consumers are charged for consuming services. Prices are updated regularly according to current aggregate supply and demand, to bring the system into equilibrium over time [14].

Trading is enabled via a centralized accounting system, where the server has the role of a bank. The server maintains an account balance for each user starting with a balance of zero and allows each user to take on a certain maximal deficit. The primary purpose of the virtual currency is to allow users to do work at different points in time while maintaining fairness. Users have a steady inflow of money from supplying resources and outflow of money from consuming services, which varies over time. In *steady state*, when the users have

⁴www.facebook.com/yoopick.

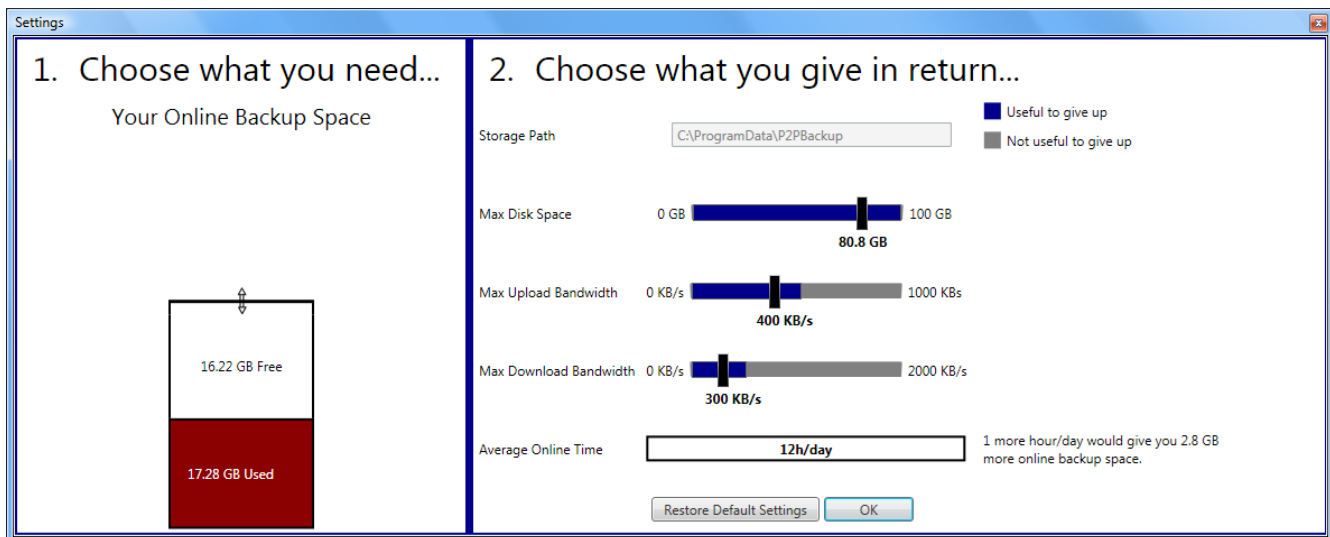


Figure 2. Screenshot of the advanced settings UI. On the left side, the user can choose the desired amount of online backup space. On the right side, the user can fine-tune the supply settings if desired. Account balances, prices and the bundle constraints are hidden from the user.

been online long enough, their income must equal their expenditure. Users cannot earn money when they are offline but must still pay for their backed up files. Thus, their balance continuously decreases during that time. As long as we do not use real money/credit cards, the maximum deficit that users can take on must be bounded. Ultimately, it is a policy decision what happens when a user hits a pre-defined deficit level. Our system will first notify the user and present options for the user to remedy the situation (e.g., increase supply), and failing this after a reasonable timeout period (e.g., 4 weeks), the user's backups will be deleted.

The server is involved in every operation, coordinating the work done by the suppliers. In the current implementation, the server allocates work to those users with the lowest account balances to drive all accounts (back) to zero over time. This is possible because the users' steady-state income must equal their expenditure. Thus, when users have been online for a sufficiently long time, their account will be close to zero. The important implication for the UI is that we do not have to show the current account balances to the users.

THE USER INTERFACE

The UI is an essential aspect of the market design because it defines the information flow between the user and the market. More precisely, different UIs invoke different mental models which lead to different user actions which can ultimately affect the market in very different ways (see Figure 1). The main questions for this UI are how the server learns about a user's individual preferences and how the user sees or "experiences" the current market prices. With primarily non-expert users and a complex, combinatorial market, direct preference elicitation methods (directly asking the users for their valuations) are infeasible to implement. However, there must be some kind of feedback loop between the market and the users for the market to have an effect. We need a hidden market that elicits users' preferences regarding their resources and communicates current prices to the users.

What You Give is What You Get

Figure 2 displays the current implementation of the UI which was built using the Windows Presentation Foundation (WPF). The user can open this "settings window" to interact with the market. This window is clearly separated into two sides: on the left side, the users can choose how much online backup space they need. On the bar chart the users can see how much they have already backed up and how much free online backup space they have left. On the right side of the window, the users can choose how much of their own resources they want to give up in return. On the top of the right side, the users see the storage path, i.e., where on their own computers the file pieces from other users are stored. Then, for each of the resources of space, upload and download bandwidth, there is a separate slider which the users can move to specify how much of that resource the system should maximally use.⁵ Below the sliders the current average online time of the users is displayed. To change this value the users have to leave their computer online for more or fewer hours per day than they are currently doing, though we can conceive of schemes in which the application can directly control such settings as power savings and hibernate mode. Next to the online time information the system also tells the users the effect of leaving their computer online for 1 more hour per day (i.e., how much more online backup space they would get in return). This is meant to make the users aware of the important role of their online time: the longer the users are online, the more useful their supply of space, upload and download bandwidth becomes, and thus the higher their income.

To change anything about their settings, the users can either drag the bar chart on the left side up or down, move any of the sliders on the right side, or change how often they are online. Both sides of the window are connected to each other such that a change on either side affects and dynamically updates the values on the other side as well. The seman-

⁵The maximum value for these sliders can be determined automatically: the limit for space is simply the free space on the users' hard drives; the bandwidth limits can be determined via speed tests.

tics of this connection are important: on average, users must pay for the total consumption chosen on the left side with the supply chosen on the right side. If the users increase any of the sliders on the right then the bar chart on the left grows because the amount of free online backup space increases. If the users decrease a slider then the bar chart on the left shrinks, because the amount of free online backup space decreases. When the user directly drags the bar chart up or down to choose how much free online backup space they want, then the three sliders on the right side move left or right, proportionally to their previous position.

The captions of the two sides of the window start with the numbers “1.” and “2.” to motivate the users to start on the left side with the supposedly easier task of choosing how much free online backup space they need, and next to move on to the right side of the window, where they can fine-tune their supply settings. Note that in practice we expect roughly two categories of users: *basic users* will only ever use the left side of the window to choose how much online backup space they need. They either do not care about which resources they give up, or they do not even understand the meaning of upload bandwidth, download bandwidth, etc. The second category of users are the *advanced users*, i.e., those users that understand the meaning and relevance of giving up their own resources and want to control their supply. In a deployed system, the settings window would initially show the left side of the window and only upon clicking an “advanced” button would the right side appear. However, note that in this work we are primarily interested in the advanced version, i.e., whether users understand the more complicated part of the UI interacting with the market. This is why in our study, the complete window is shown from the beginning.

The UI shown in Figure 2 allows users to express their preferences over consuming backup services and supplying their resources. For example, if a user needs 20 GB of free online backup space, there are several different slider settings that allow this. Some users might specify to give more space and less bandwidth, others might specify it the other way around, depending on their available resources and individual preferences. Because the user’s preferences can change over time this is not a task that can easily be automated. Note that we do not expect the users to constantly adjust their settings. Rather, we expect users to choose settings that give them enough online backup space such that they do not have to worry about their settings for a while. However, as they near their quotas, the system will notify them (via an email and visually in the application). At that point, we expect most users to adjust their sliders again, according to current market conditions and their preferences.

Combinatorial Aspects of the Market: Bundle Constraints

Now we get to the first challenge regarding the hidden market design for this application: the combinatorial nature of the market, i.e., the problem that only bundles of resources are useful to the system. In general, the free online backup space increases when the users increase one of their sliders. However, this is only true for a subset of possible slider positions. In particular, if a user keeps increasing one slider towards the maximum while the other two sliders are rel-

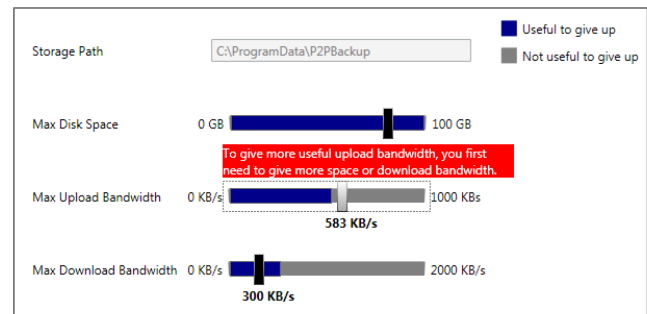


Figure 3. When users provide more of one resource than is useful to the system, then they get notified via a small pop up window.

atively low, at some point the online backup space on the left might stop increasing. For example, if users limit their upload bandwidth to 5 KB/s, then increasing their space supply from 50 GB to 100 GB should not increase their online backup space. We would simply never store 100 GB on these users’ hard disks because 5 KB/s would not be enough to have a reasonable retrieval rate for all of these file pieces. Thus, for the system to use the whole supply of 100 GB, the users would first have to increase their supply of bandwidth. An analogous argument holds true for other combinations of resources. For example, if a user wanted to give a lot of upload bandwidth but keep the supply of space low, then at some point giving more bandwidth would not be useful. Again, to make use of the download bandwidth, the system would need to store many file pieces on that user’s computer which is not possible given the current low limit on space.⁶

To provide the user with some visual information regarding how much supply of a resource is “useful to the system” given the current other slider settings, we augmented the traditional slider UI element, building the new WPF slider control shown in Figures 2 and 3. The sliders are colored blue and gray, and the legend on the top right of the window explains the color coding. In the blue region, slider movements have an effect on the online backup space because setting the slider to any position inside that region means that the system can effectively use all of the supplied resource. The gray region of the slider is the region where slider movements no longer have an effect on the user’s online backup space because giving that much of the resource is “not useful to the system,” given the other settings. Because the colors and the legend might be difficult to understand or be overlooked, we also notify the user once the slider is moved from the blue into the gray region with a small pop up message that disappears once the mouse button is released (see Figure 3).

The color-coded sliders provide the user with all the necessary information about the bundle constraints. When one slider is moved down, the blue regions on the other two sliders first stay the same and eventually decrease. Analogously, when one slider is moved up, the blue regions on the other two sliders first increase and eventually stop increasing. If a user sets the sliders in the same ratios as the system-wide usage of all resources, they are always inside the blue re-

⁶These bundle constraints only apply to space, upload and download bandwidth. For “availability” there is no minimum or maximum supply that is useful, independent of the other resources.

gions. However, requiring this exact ratio from all users is too restrictive, ignoring the system's flexibility in allocating work. Instead, we use a "slack factor" specifying how different from the system-wide usage ratios an individual user's supply can be. Thus, we give our users the freedom to supply different ratios of their resources.

Exposing/Displaying Market Prices

Because the UI gives users some freedom in choosing their resource supply, we must price the resources correctly. In our system, prices are updated daily depending on aggregate demand and supply, moving the system into equilibrium over time. Without updating prices, we might have a supply shortage for some resources. For example, many users might decide to give lots of disk space and very little bandwidth. To counteract a shortage of bandwidth, we would increase the price of bandwidth, incentivizing users to give more bandwidth instead of space. But for this mechanism to work, it is necessary that the prices are exposed to users (directly or indirectly), so they can react and change their supply settings. For example, if the price for upload bandwidth went up relative to download bandwidth, then users might benefit from increasing their upload bandwidth supply a little and in return decreasing their download bandwidth supply a lot. Due to large market effects, no individual's choice has a significant effect on the overall system. However, the market is designed such that social welfare is maximized when all users simply optimize their own choices.

Because prices are not a natural concept in a backup application, we have chosen to hide them as much as possible. A user can "experience" the relative prices indirectly by moving the sliders while observing the bar chart on the left. If the user moves a slider a little and the bar chart only changes a little, this means that the current price for that resource is relatively low. If the user moves a slider a little and the bar chart changes a lot, this means that the current price for that resource is relatively high. Note that this is one of the essential aspects of this hidden market UI: it allows us to communicate the current market prices to a user in a non-explicit way. In particular, users can be unaware of the price-based market underlying the backup system, and yet over time they will notice that for some resources they get more in return than for others. They can then choose the supply combination that is currently best given their preferences.

Market-Based vs. Non-Market-Based Approach

We know of no other P2P backup system that also uses a market. In the other systems, the ratios between the supplied resources space, upload and download bandwidth are fixed, and the same across all users. The advantage of our market-based approach is the additional freedom we give the users. Allowing them to supply different ratios of their resources increases overall economic efficiency and makes the system more attractive for every user. Note that without using a market, this freedom would not be possible, because there would be no mechanism to incentivize the users to supply the scarce resources. However, the non-market-based approach has the advantage that users do not have to worry about the bundle constraints or prices, because they are never even exposed

to these issues. It is not clear a priori how users perceive this trade-off between the burden of a more complex UI on the one side vs. the gained freedom in specifying the supply settings on the other side. This is why we explored this question in the usability study.

EXPLORATORY STUDIES

The UI as described above is the result of an iterative design process with two explicit feedback sessions in the middle and towards the end of the design phase. One important finding is that interacting with the combinatorial market is not an easy task for most users and that the details of the UI matter a lot in terms of invoking different mental models.

Cognitive Walk-Through

We first performed an exploratory study with 6 volunteers (employees at a large software company) to test the usability of our first UI design in a cognitive walk-through. At that point we received valuable feedback regarding the shortcomings of our initial UI design. First, basic users need a simpler way of specifying how much online backup space they want than using the three sliders. This led to the development of the bar chart which is easy to manipulate and automatically adjusts all three sliders, taking care of the bundle constraints. The second finding was that when interacting with the sliders, the users had difficulty understanding the market's bundle constraints and requested more help. This led to a re-design of the sliders themselves with the blue and gray regions, the addition of a legend, and the implementation of the pop up notification.

Pilot Study

Before running the usability study, we performed a pilot study with 6 new users to evaluate our second version of the UI on non-computer scientists. Again, there were two key findings from that study. First, we found that terminology really matters. To alleviate confusion, we changed the legend of the slider color coding. Before, we had used the term "Usable resources" to denote the blue region of the sliders. It turns out that changing the labels of the legend to "Useful to give up" and "Not useful to give" up made a big change regarding users' understanding of the whole system, as they were thinking about the settings in a very operational manner. This highlights the fact that different UIs invoke different mental models which lead to different user behavior in the market. The second finding was that a few users had trouble identifying how much of each resource they were currently providing, because at the time we had displayed the current value of the slider in smaller font and in parentheses to the right of the sliders, which was hard to discover. Moving the label directly under the slider and having it in bold font made a significant change in that respect as well.

USABILITY STUDY

Upon completion of the UI design phase we ran a usability study to answer the following three questions:

1. How do users perceive the trade-off between a P2P backup system vs. a server-based backup system?

2. How do users perceive the trade-off between a market-based P2P backup system vs. a non-market-based system?
3. Is our particular user interface for a P2P backup system a usable instantiation of the hidden market paradigm?

Methodology

We recruited 16 users (8 females) from the Greater Puget Sound area. All of the users had some college education and used a computer for at least 10 hours per week. The average age of our participants was about 39, ranging from 22 to 66 years old. None of the users worked for the same company, none of them were usability experts and none of them had used a P2P backup system before. All of the users understood the meaning of “backing up your files” before coming to the study, however only a few of them had used server-based online backup systems before. We recruited two different groups of users: novices and experts. Experts were screened to be users who had used P2P file-sharing software and modified the maximum bandwidth limits of their client in the last 5 years. We also ensured they had some idea about the speeds of an average home broadband connection. Novices were screened such that they did not have technical jobs, were not sophisticated enough to set-up a wireless router by themselves, and had never adjusted the maximum bandwidth limits of a P2P file-sharing client.

In this work we are particularly interested in evaluating the “advanced settings” version of the UI. Thus, our true target group of users was in fact the experts group. However, we included the novice users to make sure we identified all of the problems of the UI or the system in general that might not be found when only testing expert users. We had 8 experts and 8 novices. We ran one participant at a time with each session lasting about 1.5 hours. The users filled out a pre-study questionnaire (20 minutes), completed a series of interactive tasks using the UI (45 minutes), and then completed another post-study survey (20 minutes). We ran the software on a single 3 GHZ Dell computer at full resolution using a 20” 1600x1200 Syncmaster display.

Users’ View of P2P Backup Systems

In this section we aim to answer Question #1, i.e., how users view the trade-off between a P2P backup system and server-based backup. Each of our 22 users (6 from pilot and 16 from usability study) filled out a pre-study survey with 25 simple questions regarding the two backup alternatives. In general, we found that security and privacy are the two main concerns users have about online backup (multiple choice question with multiple answers permitted). However, the concern that the service might be too expensive already ranks third (for 11 out of 22 users), higher than concerns about the speed of the service (10/22), the reliability of the service (7/22) and the complicated use of the service (7/22). Users were definitely not willing to pay much money for backup services. Nearly all users (21/22) indicate being more likely to do backup if the service were free. But more importantly, 13 out of the 22 users would only be willing to pay \$50 or less per year for backup. Given that the current market price for server-based systems is around \$120, these users seem like good candidates for using a free P2P-based alternative.

When asked what they liked least about P2P backup systems, users generally raise concerns about security, privacy and reliability, viruses and the general disutility for giving up disk space and bandwidth. On the latter point, we obtained one interesting data point from a survey performed by a large software company in 2008, involving 28,000 users with 13,000 machines: about 40% of Windows users have more than half of their hard disk free. These users could be good candidates for using a P2P backup system, because if willing, they would be able to give up large amounts of hard disk space. The above concerns about P2P backup were perhaps alleviated with what the users liked most about our P2P system, which included that it broke files into little pieces and geographically distributed data rather than keeping it all in one place, and that it was free. Without any additional information or guarantees about the security and privacy of a P2P backup service, 12 out of 22 users said they would consider using a P2P-based backup service. But users gave an average of 3.9 (on a 5-point Likert scale, with 1=very unbelievable and 5=very believable) regarding the believability of the properties of the P2P backup system. Thus, a good number of users still had some doubts that this system can work, even after a detailed explanation. Interestingly, when asked about their opinion regarding the security and safety of a server-based vs. a P2P-based backup system, both offered by a company like Microsoft, 18 out of the 22 users gave both systems the same score on a 5-point Likert scale. The average score was only slightly worse for the P2P backup system namely 3.6, compared to 3.9.

Interactive Part: Study Set-up

The purpose of the interactive part of the usability study was to evaluate how users understand the hidden market UI, which mental models are invoked and whether users can successfully interact with the market. Note that the users interacted with the real P2P backup client software that was connected via TCP to the P2P server application and to 100 other simulated clients. We started the users off with two warm-up tasks. First, they had to perform one backup using the software. Second, they had to open the settings window and answer a series of questions regarding the information they see. Upon completing this warm-up phase, the interesting part of the interaction started.

We gave the study participants 11 tasks, each consisting of a user scenario with hypothetical preferences, and a description of the goal setting for that user. We chose tasks with varying complexity and we also tested different mental models in different tasks. For example, Scenario 1 was the most simple one, asking the user to “change the settings such that you have approximately 15 GB of free online backup space available.” In contrast, Scenario 11 was rather complex, asking the user to “imagine you are a user who likes to download videos and store them on your computer for a while. Assume that you need 20 GB of your own hard disk space to store the videos, and obviously you need lots of download bandwidth, but you do not care too much about upload bandwidth. Please change your settings so that you have approximately 25 GB of free online backup space available while taking the other constraints into account.”

Cat.	Mental Model	Experts	Novices	Total
1	Give & Take	8/8	8/8	16/16
2	Bundling	4/8	5/8	9/16
3	Prices	5/8	2/8	7/16
4	Bundling (Learned)	5/8	6/8	11/16

Table 2. Results from the Usability Study: Number of Users Falling into Comprehension Categories

We asked the users to “think out loud” as they performed each task and we made detailed observations during the tasks. Using the 11 tasks, we tested four different mental models, i.e., aspects of the user’s understanding of the market:

1. **Give & Take:** The users understand they must give some of their resources (on the right side) and get a proportional amount of online backup space in return (on the left side). This was tested using tasks 1 and 2. The test was deemed successful if the users adjusted all settings correctly.
2. **Bundling:** The users understand the bundle constraints, i.e., that they cannot provide zero of any resource because only resource bundles have value. This was tested using tasks 3 and 4. The test was deemed successful if the users adjusted all settings correctly.
3. **Prices:** The users understand that different resources can have different “prices” at different points in time. This was tested using tasks 7, 8, and 9. The test was deemed successful if the users adjusted the settings for task 9 correctly (tasks 7 and 8 gave them practice to learn the model).
4. **Bundling (Learned):** The users understand the bundle constraints after exploring the UI for a while, i.e., after a certain learning period. This was tested using tasks 10 and 11. The test was deemed successful if the users adjusted all settings correctly.

Note that the tasks were set-up such that finding the correct setting by coincidence was unlikely. The correct setting was often so unnatural that the user researcher could easily decide whether the participant had truly understood the task (and thus the right mental model had been activated) or not. Of course, the “think out loud” method also helped determining the result of a test. For example, when testing the understanding of the bundle constraints, if a user said something like “Ah, I obviously cannot give 5GB of space without giving any bandwidth, thus I choose to supply the minimum amount of bandwidth I have to give,” then this counted as sufficient understanding of the bundle constraints. The rare cases where a user had coincidentally chosen the correct settings but did not display sufficient understanding of the problem were also deemed to be failures in our experiment.

Results from the Usability Study

Table 2 summarizes the results from the usability study, evaluating whether the 4 different mental models have been successfully activated or not. It turns out that the basic aspects of the UI were understood by all users (1: Give & Take). However, the first time the users faced a combinatorial task, e.g., “minimize your upload bandwidth while maintaining at least 15 GB of free online backup space”, only 9 out of 16 users completely understood the problem and found the opti-

mal settings. The understanding of the bundle constraints of the market improved towards the end of the study, showing that a certain learning effect had occurred. In particular, 2 of the users that had not understood the bundle constraints at the beginning, understood them well at the end of the study, leading to 11/16 successful outcomes for “Bundling (Learned)”. The most difficult tasks were certainly the ones testing the users’ understanding of prices because this required three steps from them: first, discovering that different resources had different prices, second, understanding the implication for their supply of resources, and then third, choosing the optimal supply settings for themselves given current prices. Only 7 out of 16 users successfully completed all three steps, and thus were deemed to “understand the pricing aspect.” We will discuss the implications of these findings for the future market and UI design in a later section.

One interesting finding that shows up immediately is that the performance of the users is uncorrelated with the way we segmented them into experts or novices (see Table 2). Thus, prior experience with P2P file-sharing software did not really matter. Instead, anecdotal evidence suggests that those users whose jobs or education involved some mathematical modeling seemed to understand the concepts underlying the UI faster. In some sense, the tasks we gave the users can be seen as IQ test tasks. However, a factor that is difficult to measure but did play an important role in this study is the users’ curiosity, i.e., how much the users liked to play with the sliders until they figured out how the interface worked.

Post-Study Survey Results

Upon completion of the interactive part of the study we gave the users another questionnaire with 17 questions, asking them about their experience with the UI and other feedback. Despite the fact that almost every user had difficulties with at least one of the tasks, the user feedback was largely positive. Most users thought that the software made it easy to perform the tasks they were given (with a 3.8 average on a 5-point Likert scale, with 1=strongly disagree and 5=strongly agree) and they indicated that they enjoyed using the UI (3.8 average on the same 5-point Likert scale). Most users were pretty confident that they completed the tasks successfully (with an average 4.0 on the same 5-point Likert scale).

The users liked the graphical/visual representation of the concepts involved. Despite some difficulties with solving the tasks, the users thought that the UI was “clean, simple, intuitive and easy to use.” All users loved the ease of using the bar chart to choose the desired amount of free online backup space. Furthermore, they liked that the UI gave immediate feedback regarding the consequences of their choices. The users primarily disliked that it took them a while to understand the concept and logic behind the sliders.

Users’ View of “Market-Based vs. Non-Market-Based”

From the pre-study questionnaire we have seen that for a large number of users, P2P backup systems could be an attractive alternative to server-based systems. However, this still leaves open the question how users perceive the trade-off between a market-based system (that gives users more freedom in choosing different combinations of supplied re-

sources) vs. a non-market-based system (that has a simpler UI). In the post-study questionnaire we asked the users twice to compare the two options. The first time we asked the question, we gave no additional information beforehand. But before asking them for the second time, we described a particular scenario where a user of the non-market-based system could not simultaneously use the backup system and watch video-on-demand, while this would be possible with the market-based system due to the larger freedom in choosing the supplied resources. The results were that, when asked for the first time, the users already slightly preferred the market-based system (3.3 on a 5-point Likert scale, with 1=definitely prefer the simpler UI and 5=definitely prefer the complex UI). After describing the particular scenario where the non-market-based system would lead to limitations, the average score rose to 4.0. We interpret these results as follows: a priori, some users do not see the advantage of a market-based system. However, after realizing the possible limitations of the non-market-based system, they see the benefit of the increased freedom in choosing what to supply, and they value this benefit higher than the disutility from the additional complexity of the UI.

DISCUSSION

Implications for future UI Design

One clear finding from the usability study is that the separation of the settings window into two sides was a good decision, and that the left side of the window (pulling the bar chart up and down) is very usable, for all users. Thus, this can definitely serve as the “basic UI” for users who are not interested in fine-tuning the settings. When we turn attention to the right side of the window, things get more complicated. The design of the sliders, including the blue and gray regions and the pop-up window, invoked the right mental model (understanding the bundle constraints of the market) in about half of our users, but the other half struggled with this concept. We still consider this a good starting point for two reasons: first, the users who did not understand the system would likely fall into the category of “basic” users that do not even worry about fine-tuning their settings. Second, even if the users wanted to fine-tune their settings and did not understand the bundling aspect, there are no invalid settings, just suboptimal ones. As long as they keep moving the sliders until they achieve the desired amount of online backup space, they will end up with valid settings. Furthermore, the findings regarding the learning effect were positive, i.e., after playing with the UI for a while, 2 more users figured out the bundle constraints of the system. Thus, it is to be expected that in a deployed system, users would increase their understanding of the UI over time. However, our observations from the usability study also helped us re-consider the UI design to improve the users’ understanding further. Some users who were struggling suggested more information and help buttons/texts. Thus, we are currently investigating how to incorporate better user assistance into the UI.

Now we get to the most difficult part of the UI design: exposing/displaying prices. In our study, only 7 out of 16 users passed the “pricing test.” We attribute this mainly to the difficulty the users had with discovering the prices in the first

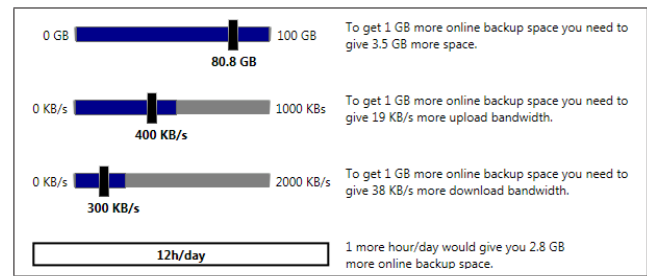


Figure 4. Future UI Proposal: surfacing the price information to the user by displaying additional help text next to the sliders.

place, and not to their conceptual understanding. To discover the fact that different resources have different prices, the users had to play with the sliders while observing how their online backup space changes. Based on our observations, once the users discovered the prices, they usually also understood the implications and were able to find the optimal settings given current prices. We are currently considering surfacing more of the price information to the users. We do not want to display \$-values, but we want to point the users more directly to the fact that for some resources they get more in return than for others. Our current proposal for how to better surface the price information can be seen in Figure 4. The only change is that to the right of each slider we show the users how much more of that resource they have to provide to get 1 additional GB of online backup space.

One of our more surprising observations was that 4 users got confused about the directionality of the sliders. Oftentimes, they used the sliders correctly for 5 or 6 tasks and then suddenly got confused. However, they realized the confusion once they started moving the slider, seeing that the bar chart moved in the wrong direction. One possible explanation is that the users are having a hard time understanding the semantics of the blue and gray regions of the slider. Even though the right side of the window has a large caption saying “Choose what you give up...”, it seems that for some users, the particular design of the sliders invoked the wrong mental model: they thought that the blue regions meant how much of the resources they keep and consequently the gray region meant how much they give up. Once the users are in that mindset, it is natural that they move the slider to the left when trying to give up more resources. Another observation was that even the users who did not misunderstand the directionality of the sliders had the strong tendency to move the sliders to the intersection of the blue and the gray regions, perhaps because that looks like an optimal setting, which it is generally not. We are now considering alternative designs for the sliders that visualize the bundle constraints in a way that is more intuitive for more users.

Implications for future Market Design

The fact that only 7 out of 16 users discovered/understood the pricing aspect of the market is also an important issue for the economic market design for two reasons: first, if individual users do not understand that different resources have different prices they might choose suboptimal settings for themselves which decreases total social welfare. But more importantly, if too many users do not understand prices, then at some point the underlying market is adversely affected.

This is where we are closing the loop: we started with an idea of a P2P resource exchange market, designed the market and the hidden market UI, tested the UI, and have now found that the choice of the UI has huge implications for the underlying market. This suggests that making adjustments to the market design itself might be necessary.

The main problem for the market is that if too many users do not react to prices then the price update algorithm might not work anymore. Remember that once a day the server looks at the aggregate demand and supply of all resources and updates market prices to move supply and demand towards a desirable equilibrium (see [14] for details regarding the equilibrium and the price update). However, the equilibrium concept as well as the price update algorithm rely on the standard economic assumption of fully rational users that make optimal choices for themselves given current prices. Fortunately, a small numbers of price-insensitive users is not a problem for the algorithm because it operates on aggregate numbers. Thus, in our ongoing work regarding the design of the underlying market we are analyzing the effect of having a significant number of users who are insensitive to price changes. Via simulations we study the behavior of the market given various kinds of user profiles. It is our goal to make the market design more robust such that a higher number of price-insensitive users can be tolerated. Based on the design of the price update algorithm we already found two ways of addressing this issue: first, we can slow down the frequency of the price updates and second, we can adjust a softening parameter internal to the algorithm.

CONCLUSION

In this paper we have explored the intersection of UI design and market design with the goal of making a contribution towards designing better UIs for complex, non-traditional markets. We have introduced a new design paradigm called “hidden markets” and tested one particular instantiation in a P2P backup domain. We have provided a detailed description of the new hidden market UI we designed and how it hides the complexities of the market. The new slider control we developed uses colored regions to visualize the combinatorial market constraints to the user in an intuitive way. The resource prices can be experienced indirectly when the users move the individual sliders while observing how their online backup space changes. In a formative usability study we have found that many users would consider using market-based P2P backup systems. We have evaluated our proposal of a hidden market UI, finding that users were quickly able to use the UI to solve basic tasks and most users understood the market’s bundle constraints. However, we also identified some shortcomings, most importantly that users had trouble discovering the fact that different resources have different prices, and thus many users remained price-insensitive. We suggested some ways to address these shortcomings by making changes to the market and the UI design. Even though there are still many open problems, this study provides the first encouraging results and shows promise for hidden markets. We hope that the general paradigm of hidden market design will be adopted more widely, paving the way for more successful market-based systems in the future.

ACKNOWLEDGEMENTS

The authors would like to thank the user study participants. We also thank Denis Charles, Max Chickering and Sidd Puri who were involved in the market design [14] and in the design of an earlier version of the UI [13]. The first author was supported by a Microsoft Research PhD Fellowship.

REFERENCES

1. C. Aperjis and R. Johari. A peer-to-peer system as an exchange economy. In *Proceedings of GameNets*, 2006.
2. S. Goel, D. Pennock, D. M. Reeves, and C. Yu. Yoopick: A combinatorial sports prediction market. In *Proceedings of AAAI*, Chicago, IL, July 2008.
3. J. Hahn. The dynamics of mass online marketplaces: a case study of an online auction. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, Seattle, WA, April 2001.
4. F. A. Hayek. The uses of knowledge in society. *American Economic Review*, 35:519–530, 1945.
5. G. Hsieh and S. Counts. mimir: A market-based real-time question and answer service. In *Proceedings of CHI*, Boston, MA, April 2009.
6. G. Hsieh, R. Kraut, S. E. Hudson, and R. Weber. Can markets help? applying market mechanisms to improve synchronous communication. In *Proc. of CSCW*, 2008.
7. K. Lai, L. Rasmusson, E. Adar, L. Zhang, and B. A. Huberman. Tycoon: An implementation of a distributed, market-based resource allocation system. *Multigent Grid Systems*, 1(3):169–182, 2005.
8. J. Li and C. Zhang. Distributed hosting of web content with erasure coding and unequal weight assignment. In *Proceedings of ICME*, Taipei, Taiwan, June 2004.
9. D. Morris, I. Simon, and S. Basu. Exposing parameters of a trained dynamic model for interactive music creation. In *Proceedings of AAAI*, Chicago, July 2008.
10. S. S. Parikh and G. L. Lohse. Electronic futures markets versus floor trading: implications for interface design. In *Proceedings of CHI*, Denver, CO, May 1995.
11. A. E. Roth. What have we learned from market design? *The Economic Journal*, 118:285–310, 2008.
12. T. Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.
13. S. Seuken, D. Charles, M. Chickering, and S. Puri. Designing User Interfaces for Hidden Markets. In *Proceedings of the IJCAI Workshop on Intelligence and Interaction*, Pasadena, CA, July 2009.
14. S. Seuken, D. Charles, M. Chickering, and S. Puri. Market Design and Analysis for a P2P Backup System. In *Proceedings of NetEcon*, Stanford, CA, July 2009.
15. I. Simon, D. Morris, and S. Basu. Mysong: automatic accompaniment generation for vocal melodies. In *Proceedings of CHI*, Florence, Italy, April 2008.