User-Directed Routing: From Theory, towards Practice

Paul Laskowski
School of Information

Benjamin Johnson
Logic and the Methodology of Science

John Chuang
School of Information

UC Berkeley

{paul,chuang}@ischool.berkeley.edu, benjamin@math.berkeley.edu

ABSTRACT

User-directed routing technologies – that is, systems in which users choose their own routes through a communications network – have generated considerable interest in recent years. Despite their numerous theoretical advantages, ISPs have so far resisted these technologies, even as users have learned to capture some routing power through overlay networks. This study responds to this disconnect between theory and practice by asking how user-directed routing would affect three prominent objectives of network operators: maintaining control over the network, earning profits, and keeping inner details of the network secret. Contrary to the modern theme in routing proposals, we argue that user-directed routing is not fundamentally incompatible with ISP-control, as long as a flexible pricing system is in place. Instead – and under surprisingly general assumptions – an ISP can use prices on the open market to induce any feasible traffic pattern. Moreover, we argue that the market-based approach maximizes welfare for any given traffic pattern. In general, our model does not guarantee whether an ISP will earn more money under user-directed routing. Nevertheless, we provide some intuition to suggest why a typical ISP may expect higher profits. Finally, we suggest that giving routing power to users conflicts with an ISP's desire for secrecy. At the same time, widespread adoption of user-directed routing, perhaps promoted through regulation, may facilitate a transparent and civil industry, to the benefit of many ISPs.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems; J.4 [**Social And Behavioral Sciences**]: Economics

General Terms

Design, Economics, Theory, Legal Aspects.

Keywords

Control, Market-Based Routing, Pricing, Traffic Engineering, Transparency, User-Directed Routing.

1. INTRODUCTION

The idea that users choose their own routes through a communications network has garnered much attention in the research community – and no wonder. *User-directed routing*, as we will call it, promises many advantages over centralized routing by a network operator. Several studies suggest that market prices may be used in such a system to maximize welfare [4][18]. Kelly further argues that for specific utility functions, market prices may be used to meet certain fairness criteria [11]. Beginning with Clark, et al., researchers have postulated that user-directed routing has the potential to improve competition and enhance service diversity [1][6][9][26]. Laskowski and Chuang argue that user-directed routing may facilitate innovation in the Internet backbone [15].

With these benefits, it is no surprise that numerous user-directed routing systems have been proposed for use on the Internet (e.g., [1][3][5][9][13][26][28][30][31][32]). As yet, however, Internet service providers (ISPs) have been reluctant to embrace these technologies, and user-directed routing in the Internet seems far from becoming reality. This disconnect between theory and practice suggests a high-level research question:

> *How would the adoption of user-directed routing affect the objectives of network operators?*

We believe that examining the perspective of ISPs is a pragmatic step towards understanding why previous routing proposals have failed, and finding ways that the routing system could evolve in the future. The need for some kind of change has never been greater than today, as the rise of overlay networks places mounting strain on an ISP's ability to implement routing policies.

Modern overlay networks, which include popular peer-to-peer file sharing applications like BitTorrent, Kazaa, and eDonkey, allow users to influence their data routes in a rather coarse way, by switching among different senders of data or by employing other users as intermediary forwarding agents. The latter technique may seem inefficient, since data must traverse a forwarding agent's local access connection twice. Despite this disadvantage, such forwarding networks have been shown to substantially improve network performance for users [3], highlighting just how different the routing preferences of users may be from those of ISPs. This misalignment of incentives has probably existed since the early days of the internet, but has only recently been made visible by overlay networks.

The upshot of these developments is that routing power is now effectively split between users and ISPs, with each attempting to pursue their interests in competing routing systems. This may be problematic for ISPs in various ways:

1. Overlay networks can make it difficult for network operators to control expenses. Peer-to-peer networks form connections without regard to an ISP's costs. As a result, over 92% of peer-to-peer traffic crosses multiple ISPs, forcing local access providers to pay extra transit fees, estimated at over €500M per year worldwide [21][22]. These fees are one reason that Norton labels Internet video "the next wave of massive disruption to the U.S. peering ecosystem" [20].

NetEcon'08, August 22, 2008, Seattle, Washington, USA.

2. Overlays can make it difficult for ISPs to design routing policies to achieve desirable network-wide traffic patterns – a task known as traffic engineering [12]. As ISPs reconfigure their networks, an overlay may respond by updating its own routing logic, possibly overriding the providers' intent. This impedes an ISP's control over its own network, making it difficult to plan future capacity, avoid congestion, and meet business obligations.
3. Complex feedback effects could occur when an ISP and an overlay (or multiple overlays) each adjust their routing logic in turn. Several studies predict that the coupled system can display oscillations [12][16][25].

These difficulties form the backdrop to our high-level research question. As we evaluate user-directed routing, we will take care to compare it against a base case in which ISPs and users each have some routing power. This comparison is not straightforward, because a typical ISP may have a great variety of objectives that will be influenced by the routing system. For this exploratory study, we will focus on just three such objectives: maintaining control over the network, earning profits, and keeping inner details of the network secret. While this list is necessarily incomplete, we have tried to include the most prominent of ISP concerns.

First, as explained by Keralapura et al. [12], ISPs want to maintain control over their own networks. As we mentioned above, overlay networks have already made it difficult for network operators to perform traffic engineering. Given that users can cause such disruption with only the limited routing power afforded by overlays, adopting a user-directed routing infrastructure may seem foolish. After all, isn't giving even more routing power to users the same as giving up more control of a network? Certainly, this belief is widely reflected in recently proposed routing protocols, which try to compromise between user-choice and ISP-control. In the typical approach, users are given a controlled degree of routing flexibility, but ISPs have the ultimate power over what options to expose (examples of this approach include [9][13][28][30][31][32]).

Such proposals aside, we will approach this issue anew, and ask exactly how much control an ISP must sacrifice under a user-directed routing framework. In fact, this study will float the idea that giving users complete routing power may actually *enhance* ISP-control. As evidence, we will present a stylized model of a network routing system. We will argue – under surprisingly general assumptions – that when users choose their own routes, an ISP can use prices on the open market to replicate any traffic pattern achievable through traditional traffic engineering (even in an idealized scenario without overlay networks).

A second objective shared by most network operators is earning money. As we mentioned, overlay networks have made it harder for ISPs to maximize profits, by directing traffic over expensive transit connections. Just as with the issue of control, it may seem natural to assume that giving users even more routing power would further erode an ISP's profits.

We will try to assess this conclusion with the help of a stylized model of consumers. First, we will show that a user-directed routing system maximizes *social welfare* for any given traffic arrangement. Less promisingly, our model does not guarantee whether ISPs will earn more money with user-directed routing. Nevertheless, we will give intuitive reasons for why a *typical* operator may expect higher profits in such a system.

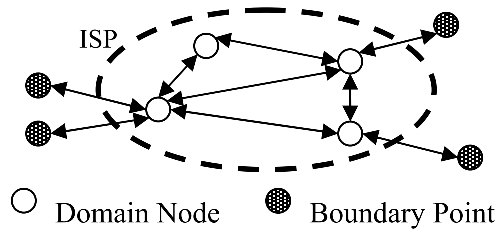

Figure 1: Network Model

The final objective we will consider relates to transparency. Operators prefer to keep the inner workings of their networks a secret. Such information could be used by competitors to make better business decisions, and possibly even target an ISP's vulnerabilities. User-directed routing would require an ISP to expose much of its structure, destroying its *informational advantage*.

We will explore the benefits that secrecy affords an ISP and confirm that most ISPs would be reluctant to embrace greater transparency. While this is a significant obstacle, we will suggest that coordinating the adoption of user-directed routing among many ISPs may yet lower the barrier to deployment.

2. CONTROL

When users choose their own routes, managing a network may seem like a daunting task. Instead of programming routers, an ISP can only adjust prices up and down to encourage users to change their behavior – a process we may call *market-based routing*. Then again, setting prices for end-users bears at least a passing resemblance to the way ISPs set link weights as inputs to popular routing algorithms. How do these two tasks actually compare?

In this section, we will provide part of the answer to that question. Our argument relies on a stylized network model, but we have aimed to keep our assumptions as minimal as possible. More striking than the assumptions we make, are the traditional ones we have left out: no utility functions, no linear or quasi-linear preferences, no convexity or monotonicity of demand. Users are only modeled in terms of their observable impact on traffic, and our main assumption is just that traffic shifts in some smooth manner as prices are varied. Because of the generality of this framework, we hope that even model-skeptics will find relevance in our conclusions.

Imagine a directed graph representing an ISP. A set of *domain nodes*, I, represents the routers within the ISP's domain, as seen in Figure 1. The remaining nodes, which we call *boundary points*, represent connections into and out of the ISP to users and adjacent providers. Each boundary point has degree one and only connects to a single domain node.

A *route* through this network is represented by an acyclic path that begins and ends at boundary points. We define R to be the set of network routes, and we'll assume our network has exactly n routes, $r_1,\dots,r_n$. A traffic pattern t assigns a real number t_i (say, in gigabits per second) to each route r_i. We may construe t as a point in $\mathbb{R}^n$.

We'll refer to each directed edge or vertex in the graph as a *component*, and we'll assume that every component has some finite capacity. If component e has capacity c, it imposes the following

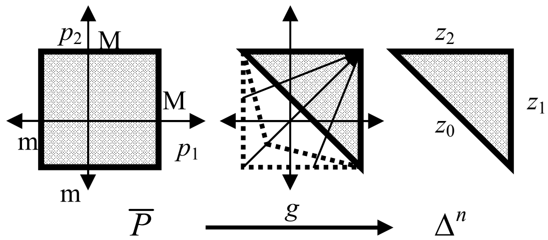

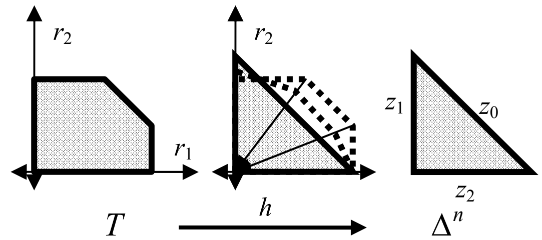

constraint on traffic: $\sum_{r_i \in R:\, e \text{ is a component in } r_i} t_i \le c$. We say that a traffic pattern t is *feasible* if it satisfies all capacity constraints.[1] Let T be the set of all feasible traffic patterns. T is a convex and bounded subset of $\mathbb{R}^n$.

Next, a pricing schema p assigns a price p_i to each route r_i. We may construe p as a point in $\mathbb{R}^n$.

Consumers are captured by a demand function D mapping price schemas P to feasible traffic patterns $T, D: P \to T$. Our main result requires three assumptions about consumers:

Assumption 1. D is continuous.[2]

Assumption 2. There is a price high enough to guarantee that no traffic flows on any route with that price: There exists M such that $p_i \ge M$ implies $D(p)_i = 0$.

Assumption 3. There is a price small enough to guarantee that any route with that price fills to capacity: There exists m such that $p_i \le m$ implies at least one component of r_i is at maximum capacity.[3]

Theorem 1. Under assumptions 1-3, the demand function $D: P \to T$ is surjective. I.e., every feasible traffic pattern can be induced through some set of route prices.

Proof sketch. We begin with a high level description of how the proof works. Based on our assumptions, we can restrict our attention to sets of prices and traffic patterns that form bounded and convex subsets of $\mathbb{R}^n$. Our assumptions also ensure that the demand function takes the boundary of the prices to the boundary of the traffic patterns. Since the demand function is continuous, its image cannot have any holes in it. Thus each of the traffic patterns must be obtainable from the set of prices.

Our proof is non-constructive in the sense that it does not tell us how to find a set of prices to generate a given traffic pattern. Instead, it uses well-developed tools from algebraic topology to derive the result by contradiction.

Let $\overline{P}$ be the set of price schemas with $m \le p_i \le M$ for each i. Let Δ^n be an n-simplex whose boundary Z has $n+1$ faces $z_0, \ldots, z_n$. Define $g: \overline{P} \to \Delta^n$ to be any homeomorphism that maps the face of $\overline{P}$ with $p_i = M$ to face z_i (for $i = 1, \ldots, n$), and maps all other faces of $\overline{P}$ to z_0. (Loosely speaking, g squashes a "cube" into a "pyramid" in a way that preserves the faces we know the most about, and combines the other faces into one, as depicted in two dimensions above.) Similarly define $h: T \to \Delta^n$ to be a homeomorphism that maps the face of T with $p_i = M$ to z_i, (for $i = 1, \ldots, n$), and all other faces of T to z_0. Assumptions 2 and 3 then imply that the induced function, $\overline{D} = hDg^{-1}$, maps each face of Δ^n to itself. If we let $\overline{D}_Z : Z \to Z$ be the restriction of $\overline{D}$ to the boundary of the simplex, $\overline{D}_Z$ is homotopic to the identity map through a straight-line homotopy. Therefore, $\overline{D}_Z$ has degree 1.

Now suppose $\overline{D}$ is not surjective, and let v be some point not in its image. Since $\overline{D}_Z$ has degree 1, its image must be all of Z; and so v must be an interior point of Δ^n. We can now use v to define a projection $\phi: \Delta^n \to \Delta^n$ that sends each point along the ray from v until it intersects Z. Since ϕ fixes Z, the composition $\phi \circ \overline{D}: \Delta^n \to Z$ is an extension of $\overline{D}_Z$ over the entire simplex. But this implies that $\overline{D}_Z$ is homotopic to the constant map, and thus has degree zero, giving us a contradiction. □

Theorem 1 suggests that market-based routing can achieve every traffic pattern that traditional traffic engineering can. Of course, knowing that prices exist to induce a traffic pattern is only half of the story; to exert control, an ISP has to actually find those prices. Admittedly, this is not a straightforward task: when the levers that an ISP pulls are prices, it is hard to predict what effect each adjustment will have. But the same thing can be said of traditional traffic engineering, especially when customers exert competing control through overlays. In fact, by relinquishing all routing power to users, the market-based approach would create a unified system, stripping away a layer of complexity and reducing volatility. The end result may be that traffic becomes more predictable, giving ISPs more effective control, not less.[4]

[1] Constraints of this type can be used to model a broad range of concerns, including physical bottlenecks both inside and outside the network, as well as any other limitation on traffic that an ISP might care about.

[2] Continuity may be unrealistic if large customers abruptly switch all their traffic to new routes. If congestion effects are present, however, we would expect such shifts to result in suboptimal service, so large customers have an incentive to shift traffic gradually.

[3] We may allow m to be negative, representing the case that customers must be paid to fill a route to capacity. This may cause some worry, but we expect negative prices to be fully unnecessary for an ISP in practice.

[4] Due to the amount of information that must be propagated, user-directed routing at the granularity of individual links is widely

We have seen that an ISP can achieve any feasible traffic pattern by leveraging a sufficiently detailed pricing system. Of course, no such system is in place today, and creating one remains a significant obstacle to the adoption of user-directed routing. Particular attention must be paid to the user experience. Choosing routes is a complex and tedious task, and we do not expect individuals to want to choose every path by hand (except for a handful of power-users). Rather, a system must be in place to choose paths for users and present them with a simple interface. One candidate is the ROSE system presented by Lakshminarayanan, et al., in which a set of routing firms compete to advise users on good paths through the network [13]. In another possible scenario, overlay networks could take on a central routing role, purchasing their routes from ISPs and charging users a simple fee. In a well-designed system, many such overlays could coexist and compete with each other, ensuring a high-quality product. Overlays could even specialize to offer paths to particular applications, or to other overlays. Comparing these alternate scenarios remains an important open research task.

3. PROFIT

Aside from allowing ISPs to maintain control over their networks, market-based routing presents another distinct advantage. Users know what their preferences are much better than ISPs do. Intuitively, by allowing them to sort themselves into the best routes, an ISP may allocate its resources more efficiently, thereby increasing the value of its network.

In fact, we can make this more concrete: Our first theorem suggested that any feasible traffic pattern can be induced through route prices on the open market. Under modest assumptions, we can show that such a solution maximizes welfare over all possible assignments of users to routes that match that traffic pattern. Holding the aggregate traffic constant, no manual rearrangement of users by an ISP can improve the value of the network.

Let us assume that demand comes from a set of rational consumers with quasilinear utility, each wishing to send a set of flows. If a flow gets sent, its owner enjoys non-negative utility which depends on the route taken, and possibly the total network traffic – independent of how individual flows combine to yield the global traffic pattern. We will also assume that each individual's contribution to total traffic is infinitesimal. (This assumption allows each consumer to ignore the impact of her own traffic when selecting a route). Under these assumptions, we have,

> **Theorem 2.** Over all possible ways to arrange flows to match a given traffic pattern, any arrangement induced by market prices in equilibrium maximizes welfare.

Proof. Let X be an equilibrium induced by market prices, and let Y be some other arrangement of flows that yields the same traffic on each route. Let Z be the hypothetical arrangement in which each customer is given the same route she would have in Y, but is charged the price assigned to that route in X. Welfare is insensitive to money transfers, so welfare in Z is the same as in Y. Since X is at market equilibrium with respect to its route configuration, traffic pattern, and price schema, all of which are shared by Z, each consumer's utility in X must be at least as high as in Z. (If not, the rational consumer, who can safely ignore her own infinitesimal contribution to traffic in X and Z, would switch to the route in Z yielding higher utility). In addition, the ISP's profit in Z is the same as in X because the traffic patterns and pricing are both the same. Thus the total welfare in X (which is the sum of consumer surplus and ISP profit) is at least as high as the welfare in Z, which is the welfare in Y. □

Taken together, Theorems 1 and 2 imply that any feasible traffic pattern can be achieved through market-based routing, yielding at least as much welfare as any other method. This applies to all feasible traffic patterns, including the one generated by the socially optimal assignment of users to routes (if such an optimum exists). Our theorems then imply that route prices can be found that maximize the total welfare of the system. This mirrors a classic result from network models based on linear programming (e.g. [18][11]). In contrast to these studies, our approach relaxes the assumption of convexity to mere continuity, but adds the assumption of infinitesimal users.

While market-based routing enhances welfare, we cannot say for certain where the extra value goes. In particular, there is no universal guarantee that an ISP will earn more profit through a market-based approach. We can, however, offer some initial observations to suggest why market-based routing may increase profit for a *typical* ISP.

1. Under the market-based approach, an ISP can tailor routes for niche communities with specific requirements, distinguishing itself from competitors and charging a premium. Such niche communities are an untapped resource that may grow more profitable in the future as new applications are created to take advantage of tailored routes. Market-based routing also allows an ISP to exploit the diversity within its network, by matching routes with unique characteristics to users that will pay the most for them.
2. Market-based routing would allow ISPs to market innovations directly to end-users, enhancing returns on investment. This may be particularly attractive to backbone providers, who currently face a highly commoditized market that fails to reward innovation (Laskowski and Chuang explore structural factors that contribute to this [14][15]).
3. A market-based approach allows ISPs to tailor the price of routes to their underlying costs. This is especially relevant to access networks, who could set prices to recoup the expense of peer-to-peer traffic sent over transit connections. Over time, we may expect consumers to put pressure on peer-to-peer networks to adapt and avoid paths with high prices. Thus, a new generation of overlays may emerge that is sensitive to both quality and cost.

4. INFORMATIONAL ADVANTAGE

Aside from control and profit, an ISP will note that user-directed routing dramatically increases transparency, allowing users – and competitors – to observe the network's operation on an unprecedented scale. In the long-term, this effect holds important consequences for an ISP's strategic goals, and for the broader health of the network industry.

As the network neutrality debate has highlighted, there are many situations in which an ISP may benefit by discriminating against particular traffic flows. This was famously witnessed in 2004 when

considered impractical. Researchers often favor a hybrid approach in which users select an AS-level path, and each ISP determines routes within its domain. We have focused on the more radical version of user-directed routing because that is the difficult case. Our approach is equally effective when applied to the hybrid system. Simply represent the ISP as a single interior node, and Theorem 1 implies that prices can induce any feasible traffic rates between pairs of boundary points.

telecom firm Madison River Communications blocked Vonage's voice over IP (VoIP) traffic on their networks [17]. Its motivation was clear: VoIP threatened to draw customers away from their traditional phone service. In another episode, Canadian provider Telus blocked access in 2005 to a website that portrayed the firm negatively during a labor dispute [29]. More recently, network operator Comcast attracted widespread scrutiny for disrupting the popular BitTorrent file-sharing application [8].

Unfortunately, we only know about the misbehavior of these ISPs because of their brazen approach – completely severing a communication is a highly visible act. Chances are that many more episodes of discrimination go undetected. This is because traffic engineering allows an ISP to discriminate against specific traffic much more discreetly, by sending it over poor quality routes. Peha lists a variety of damaging ways that an ISP might use such discrimination, from protecting legacy services by stifling competing innovations to cornering network services [23]. Equipment vendors even offer advice on how ISPs can use discrimination to extract rents from upstream markets [23].

Providers may also apply traffic engineering to target other ISPs. In a fascinating survey, Norton lists a variety of ways that an ISP can manipulate traffic to target competitors and rise through the ISP hierarchy [19]. Some of these are described by competitors as "evil, clever, and anti-social." At times, ISPs have even been known to use "fake" traffic from a traffic generator or web spiders [19]!

Because centralized routing by ISPs allows them to operate in relative secrecy, abuses of power are hard to detect and even harder to confront. Traffic flows may receive different treatment for a variety of benign reasons, so even if a user knows that her route is poor, it is hard to establish that this is a part of a deliberate strategy. Moreover, because there are no route prices to reference, the user will find it hard to argue about the value she is missing out on.

Given the current network environment, any single ISP that adopts user-directed routing is likely to find that the resulting increase in transparency presents a variety of disadvantages. The provider may have to abandon discriminatory behaviors that it had been using to its benefit. Meanwhile, competitors will be able to observe the ISP's interior structure, and use this information to help plan future capacity, draw away customers, or negotiate advantageous contract terms. The ISP may even find itself increasingly vulnerable to the traffic manipulation tactics of its competitors. That may be enough to dissuade even the most enthusiastic technology adopter.

Of course, even if individual providers prefer to keep their internal workings hidden, the industry as a whole may benefit from a collective move towards more transparency. Widespread secrecy on the part of ISPs contributes to a murky and complex business environment, imposing a variety of costs over time. These include inefficient investment stemming from poor knowledge of the market, the need to adapt to a variety of hostile threats, and lost competition due to monopoly-seeking and informal collusion. A tangled industry is also difficult to regulate, and can serve as an incubator for increasingly anticompetitive behaviors. More importantly for users, ISPs may interfere with new technologies, degrading the rate of innovation. This affects the development trajectory of the network and network applications over the long term.

A widespread move towards transparency could therefore benefit a variety of stakeholders, including many ISPs. These providers would be willing to suffer the loss of secrecy associated with user-directed routing, if only other ISPs were required to do likewise. We therefore speculate that coordinated action, perhaps facilitated by regulation, may serve to lower the barriers to adopt user-directed routing.

5. CONCLUSION

By examining the objectives of ISPs, our analysis sheds light on the major obstacles preventing the adoption of user-directed routing in the Internet. It has long been argued that ISPs resist such technologies because they undermine a provider's ability to control its own network. This helps explain why past routing proposals have failed, but our analysis suggests that this argument does not apply to user-directed routing in general. Contrary to the prevailing wisdom, our model predicts that there is no fundamental incompatibility between user-directed routing and ISP-control. Instead, an ISP can induce any feasible traffic pattern under user-directed routing – as long as the payment system is flexible enough. Previous routing proposals all lacked this critical ingredient for success.

Developing a payment system for use with user-directed routing entails many open research questions. Who should select paths on behalf of users? We have seen that operators require flexible prices to maintain control of their networks, but how much flexibility is enough? Our model envisions attaching prices to individual routes, but the list of such routes is enormous. It may be technically infeasible to disseminate all of these prices across the network to users. Because of this, we would like to know how much control is lost under a simpler pricing format. For example, can any reasonable traffic pattern be *approximated* by attaching prices to individual links or switches? A parallel may be found in today's traffic engineering system. Instead of exercising all of its routing power, a typical ISP will prefer to use a minimal spanning tree algorithm for routing within its domain. With such an algorithm, routing becomes a much simpler matter of setting a logical weight for each link, and this benefit outweighs the resulting loss of fine-grained traffic control.

Of course, before deploying a new payment system, ISPs will be curious to know how large the new payments will be. Our model provides no universal guarantee that ISPs will earn more under user-directed routing. Nevertheless, we gave reasons for why a typical ISP might expect higher profits under such a system. We therefore expect user-directed routing to be quite compatible with an ISP's objective of maximizing profits.

The same cannot be said for an ISP's desire to maintain secrecy. We have observed that user-directed routing allows competitors to observe an ISP's inner workings on an unprecedented scale, creating a variety of disadvantages. Our analysis suggests that ISPs may partially mitigate these disadvantages by coordinating the adoption of user-directed routing with each other. Of course, regulation would be a possible way to facilitate this process, but we will defer such considerations to future studies.

Although our study has focused on how user-directed routing would affect ISPs, we were fortunate to stumble upon ways that user-directed routing can benefit society in general. Our model predicts that, over all the ways of achieving a given traffic pattern, a market-based approach maximizes social welfare. We have also argued that user-directed routing can promote a civil and transparent market, preventing a variety of anticompetitive behaviors and encouraging innovation. We hope that these benefits encourage further research into user-directed routing.

6. ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers, Jens Grossklags, Scott Shenker, and Bill Woodcock for their comments. This work is supported in part by the National Science Foundation under ITR award ANI-0331659.

7. REFERENCES

[1] Afergan, M. & Wroclawski, J. On the Benefits and Feasibility of Incentive Based Routing Infrastructure. PINS 2004.

[2] Aggarwal, V., Feldmann, A. & Scheideler, C. Can ISPs and P2P Users Cooperate for Improved Performance? SIGCOMM Comput. Commun. Rev., 37(3):29–40, 2007.

[3] Andersen, D., Balakrishnan, H., Kaashoek, F. & Morris, R. Resilient overlay networks. ACM SOSP 2001.

[4] Bell, C. E., S. Stidham. Individual versus social optimization in the allocation of customers to alternative servers. *Management Sci.* 29: 831–839, 1983.

[5] Chu, Y.-H., Rao, S. G., & Zhang, H. A Case for End System Multicast. *ACM SIGMETRICS* (2000).

[6] Clark, D. D., Wroclawski, J., Sollins, K. R. & Braden, R. Tussle in cyberspace: Defining tomorrow's internet. ACM SIGCOMM 2002.

[7] Computer Science and Telecommunications Board, National Research Council. *Looking Over the Fence at Networks: A Neighbor's View of Networking Research*. The National Academies Press. 2001.

[8] Eckersley, P., von Lohmann, F., Schoen, A. Packet Forgery By ISPs: A Report on the Comcast Affair. Electronic Frontier Foundation. 2007.

[9] Estan, C., Akella, A. & Banerjee, S. Achieving Good End-to-End Service Using Bill-Pay. HotNets 2006.

[10] He, J., Bresler, M., Chiang, M. & Rexford, J. Towards Multi-Layer Traffic Engineering: Optimization of Congestion Control and Routing. *IEEE Journal on Selected Areas in Communications* (June 2007).

[11] Kelly, F. P. Charging and Rate Control for Elastic Traffic. *European Transactions on Telecommunications*, 8: 33--37, 1997.

[12] Keralapura, R., Taft, N., Chuah, C., Iannaccone, G. Can ISPs Take the Heat from Overlay Networks? HOTNETS 2004.

[13] Lakshminarayanan, K., Stoica, I. & Shenker, S. Routing as a Service. UC Berkeley EECS Technical Report No. UCB/CSD-04-1327 2004.

[14] Laskowski, P. & Chuang, J. Network Monitors and Contracting Systems: Competition and Innovation. ACM SIGCOMM 2006.

[15] Laskowski, P. & Chuang, J. The Interplay of Routing and Pricing in the Next Generation Information Infrastructure. Under submission., 2008.

[16] Liu, Y., Zhang, H., Gong, W. & Towsley, D. On the Interaction Between Overlay Routing and Traffic Engineering. IEEE INFOCOM 2005.

[17] Madison River Communications LLC, Consent Decree, 20 FCC Rcd 4,295 (2005).

[18] Masuda, Y. & Whang, S.J. Capacity Management in Decentralized Networks. *Management Sci.* 48(12): 1628-1634, 2002.

[19] Norton, W. The Art of Peering: The Peering Playbook. Equinix Inc. 2002.

[20] Norton, W. Video Internet: The Next Wave of Disruption to the US Peering Ecosystem. NANOG 40, 2007.

[21] Parker, A. The True Picture of Peer-to-Peer Filesharing. CacheLogic 2004.

[22] Parker, A. P2P in 2005. CacheLogic 2005.

[23] Peha, J. The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy. *International Journal of Communication*, 2007.

[24] Peterson, L., Shenker, S. & Turner, J. Overcoming the Internet Impasse through Virtualization. HotNets 2004.

[25] Qiu, L., Yang, Y. R., Zhang, Y. & Shenker, S. On selfish routing in internet-like environments. ACM SIGCOMM 2003.

[26] Ratnasamy, S., Shenker, S. & McCanne, S. Towards an Evolvable Internet Architecture. SIGCOMM 2005.

[27] Roughgarden, T. & Tardos, E. How bad is selfish routing? Journal of the ACM (JACM), 2002.

[28] Snoeren, A. & Raghavan, B. Decoupling policy from mechanism in internet routing. SIGCOMM Comput. Commun. Rev., 34(1):81–86, 2004.

[29] "Telus cuts subscriber access to pro-union website", CBC News, 2005. Retrieved on 2007-08-02. http://www.cbc.ca/canada/story/2005/07/24/telus-sites050724.html

[30] Xu, W. & Rexford, J. MIRO: Multi-path Interdomain ROuting. SIGCOMM 2006.

[31] Yang, X. NIRA: A New Internet Routing Architecture. SIGCOMM Comp. Comm. Rev. 33(4):301–312, 2003.

[32] Yang, X. & Wetherall, D. Source Selectable Path Diversity via Route Deflections. SIGCOMM 2006.