# On the Benefits and Feasibility of Incentive Based Routing Infrastructure

Mike Afergan
afergan@mit.edu

John Wroclawski
jtw@mit.edu

## Abstract

Routing on the Internet today is as much about money as it is traffic. The business relationships of an ISP largely dictate its routing policy and drive the work of its engineers. In today's routing mechanism, this leads to a number of well-known pathologies. This structure is further challenged by the emergence of user-directed routing.

This paper explores these challenges and argues for the *introduction of explicit incentives (prices) into the routing fabric of the Internet.* We argue that doing so addresses limitations of the current system that are significant today and will only be exacerbated by user-directed routing. To support this claim, we describe the benefits and properties of incentive-based routing frameworks and demonstrate how such frameworks can be applied to a number of routing architectures, including BGP.

**Categories and Subject Descriptors:** C.2.2 [Network Protocols]: Routing Protocols

**General Terms:** Design, Economics, Management.

## 1. INTRODUCTION

Several attributes are desirable in Internet-scale inter-domain routing systems. Critical among these are scalability and decentralization. Today, this is achieved through a distributed system in which communication between networks happens pair-wise. Each network makes its decisions locally and transmits a limited view of the Internet onto its neighbors. It is therefore of vital importance that a given network be able to make efficient decisions myopically – based solely on the local information that it has at hand. Without this property, networks, and thus the Internet, would constantly be in flux.

This ability to make a decision without having to reason about the strategies and details of the other players is essentially that of a (weakly) dominant strategy. A model of inter-domain routing proposed by Gao and Rexford [1] argues that under a certain set of assumptions the structure of inter-ISP business relationships induces an equilibrium where each AS acts based solely on its contracts but "where no AS would change their routes," a notion they call "stability." This notion of stability can be cast as the existence of a dominant strategy in a game of incomplete information where each player's (private) type is its set of business relationships. To obtain their result, Gao and Rexford implicitly assume that traffic patterns are under the control of ISPs and that this control allows inter-ISP business relationships to be relatively stable.

In this paper, we argue that these assumptions are both difficult to maintain in the current Internet and under increasing pressure from emerging user-directed routing technologies. We suggest that user-directed routing exacerbates underlying problems in the current system by creating a significant misalignment of incentives. As a result, the routing problem is no longer solvable in dominant strategies and instead requires ISPs to predict future traffic patterns and guess how their neighbors will react. These factors place ISPs and their users (both individuals and other ISPs) in direct conflict, an example of the tension Clark *et al* refer to as a "tussle." [2]

We consider how to resolve this tussle. Hypothesizing that explicit representation and manipulation of incentives within the routing protocols is the right approach; we explore possible benefits and discuss ways in which such mechanisms could be used. Instead of creating a conflict between the networks and those seeking to influence routing, as the current system does, we explicitly bring these users into the game through incentives. We suggest that this approach will both reduce the high overheads of today's system in the current environment and respond to the new demands of the evolving user-directed routing environment in a way that today's system cannot.

The contributions of this paper are:

1. Analysis of the advantages and disadvantages of the coupling (or lack thereof) between the Internet's underlying business relationships and technical routing protocols.

2. Delineation of the core conflict between user-choice routing and current routing practices and business relationships.

3. The observation that both the current disadvantages and the now-developing conflict can be mitigated by a single, simple idea: the introduction of explicitly represented incentives (as prices) into the routing fabric.

4. Application of the approach to example routing frameworks, showing that simple changes can address the key conflict and *simplify* the technical routing environment.

5. Introduction of "The Price of Tyranny," a framework for representing and analyzing the mechanisms discussed.

6. Presentation of several research questions, drawing on networking and mechanism design, that seek to understand the implications of these mechanisms and remaining issues.


*SIGCOMM'04 Workshops,* Aug. 30+Sept. 3, 2004, Portland, Oregon, USA.

## 2. THE NATURE OF TODAY'S INTERNET ROUTING

In theory, today's Internet routing system allows ISPs to operate with arbitrarily complex and independent policies. In practice, however, ISP policy is normally quite simple and driven by one motivating factor – money.[1] The policy of these ISPs can be characterized as an exercise in cost minimization: *"Given that I must forward this packet, what is the cheapest route on which to do so?"*[2]

These business relationships fall into two broad classes. The first is the customer-provider relationship. Here one network pays the other for the traffic that passes over the link between these two networks. Details of this pricing may vary, but in each case the overall agreement is defined by the customer-provider contract. The second is peering. Peering relationships are typically formed by larger ISPs, where both agree that for traffic on routes advertised to each other, no payment will be required. While there are exceptions, peering relationships are typically formed when the traffic exchanged (and thus the potential money exchanged) is roughly equal.

## 3. DOWNSIDES TO THE CURRENT MODEL

There are also non-trivial downsides to the current model. At the heart of these problems is that ISP economics and users' desires are the fundamental quantities of the system, but are not represented in the routing protocols themselves. Instead, the ISPs must fall back on imprecise and indirect BGP techniques to convey behavioral incentives. As a result, both users and ISPs suffer in many ways:

1. **Sub-optimal Routing** Because the AS is the player who makes the routing decisions on behalf of the users, there exists the potential for a significant *moral hazard*. That is, given a fixed amount of traffic to route, the AS can, will, and (as a profit-maximizing firm) *should* make decisions which will decrease (or perhaps even minimize) its cost, at the expense of poorer service to the user. [3]

   A well-understood example of this phenomenon is "hot potato routing" [3], where an ISP hands packets off to peer ISPs whenever possible. This occurs because the peer appears free, a high incentive to use a possibly sub-optimal route. A second example is traffic routed to stabilize the ratio at a peering point - here the mis-incentive is the possible loss of the peering relationship if traffic becomes unbalanced. Sub-optimal routes due to such mis-incentives can cause decreased performance for end-to-end flows and BGP itself. While economics is not the only reason, it has been observed that 30-55% of the paths on the Internet are sub-optimally routed.[4]

2. **The Costs of Inter-Domain Traffic Engineering** Another cost of the current system is the work and risk associated with traffic engineering. As IP service commoditizes and profit margins decrease, ISPs and researchers have begun paying closer attention to costs of implementing policy. This process is often complex and/or manual [5]. This is costly in two ways. First, complex manual process is a financial burden to the ISPs – the process of cost-minimization through traffic engineering is itself ironically costly. Second, the complicated process can easily introduce significant errors into the routing system [5]. Simplifying and automating this particular process can both reduce costs and improve the level of service.

3. **Instability of Peering Relationships** Because peering contracts provide tremendous cost savings, ISPs exhibit *adverse selection* to obtain and maintain them. It becomes rational to take varied questionable steps to obtain and/or maintain these relationships – such as making side-deals with other ASes or entities to force traffic through an inter-connection point. A key reason that so much effort is invested in this process is that there is no graceful transition between the relationship of peer and that of customer-provider.

4. **Lack of Price Discrimination** In addition to increasing costs for the networks, the current inflexible scheme decreases potential ISP revenue. Unlike many more established systems (telephone, postal, airlines) most Internet pricing is based on a single rate applied to all usage. This pricing has the merits of simplicity and small-customer acceptance, but it is well known to reduce economic efficiency. As ISPs focus more on rates of financial return, the ability to discriminate on customer willingness to pay becomes a more important tool. Indeed, finer-grain differentiation has emerged in the maturation process of other networks such as transportation [6].

We observe that to the extent each of these problems is addressed today, it is done without explicit protocol visibility of ISP objectives and incentives. In some cases the objective is undefined; there is no standard inter-ISP quality metric. In other cases the incentive is defined but outside the reach of the decision-making protocol; inter-ISP financial incentives are defined by paper contracts, not the routing system.

Absent this information, ASes are forced to resort to complicated and imprecise tools (e.g., AS path padding, BGP communities, or any of the other sundry BGP options). Thus, no AS understands how its decisions impact its neighbors and users, nor can it communicate the cost of such decisions in a way to be effectively compensated for them. Devoid of a more efficient means of achieving an efficient equilibrium, the AS is often left to make arbitrary decisions, and to implement them in a costly process of trial and error (and consultation of MRTG graphs [7]). This process is costly, complicated, and inefficient for all parties involved.

From these observations we draw two conclusions:

***Conclusion 1:*** *Maintaining business relationships and control over traffic in the Internet today is costly to ISPs.*

***Conclusion 2:*** *Much of this cost and complexity stems from the fact that the financial incentives are not explicitly communicated in the tools and protocols that are used.*

In other words, the apparent simplicity of today's model is specious. The complexity we have removed from BGP has only created work and complexity elsewhere in the system.[4]

[1] In this paper we ignore non-commercial networks, such as government networks, where factors such as privacy may be tantamount.

[2] We ignore excessively parsimonious ISPs that simply drop large amounts of traffic, as ultimately no one will do business with them.

[3] We discuss in the next section that examples such as CDNs validate that these problems are both significant to users and addressable.

[4] By creating work, this added complexity not only decreases

# 4. USER-DIRECTED ROUTING

Due in part to the problems presented in the last section and their inimical side-effects, numerous methods to provide some element of user choice in routing have been proposed, built, and (sometimes) deployed. We refer in this paper to *user-directed routing technologies* when focusing on the broad principle rather than any particular implementation.

Today's most common examples of user-based routing are overlay networks. In industry [8] and academia [9], overlay technologies have been used to increase the reliability and performance of Internet flows. As such they exploit two fundamental facts. First is that BGP has no true notion of QoS and certainly no notion of end-to-end QoS. Therefore, inter-ISP BGP routing proceeds without this consideration. Second, overlays solve the moral hazard problem. By distributing the choice to the end-user, the only agent who is properly incented to pick the optimal route, they shift the balance of control.

Overlay networks are limited in that user routing choice is constrained by the location of the overlay nodes. Over time, there have been several proposals (e.g. [10] [11] [12]) that provide for even greater user control in route selection. A related form of user choice can be found in peer-to-peer networks, where a given resource may be found in multiple places, with the choice made by the application or P2P algorithm.

Although these user-directed routing proposals contain both common elements and sweeping differences, a key detail is that many – most notably overlays – can be created without the support of ISPs.[5] Based on this, we therefore make two observations:

**Observation 1:** User Choice increases the fluidity of traffic patterns by several orders of magnitude. If we assume that a traffic source (e.g., a large company or web site) might change its ISP once a year, but an overlay network can shift the site's traffic on the order of every 5 minutes, we are dealing with a factor of 100,000 change (five orders of magnitude). Of course, most overlays can adjust a significant fraction of traffic even faster than that. Moreover, these changes need not be stable – traffic could be rapidly shifted back and forth. These shifts significantly perturb capacity planning and peering relationships.

**Observation 2:** It is not acceptable to assume away the effects of user-directed technologies, deployed with or without the cooperation of ISPs. Akamai alone today carries about 15 percent of the web's traffic.[13] We must examine the impact on the current incentive structure on user-directed routing, and should go beyond this to examine what framework is best suited to support its growth.

# 5. THE CORE TUSSLE

From a simple understanding of Internet economics and user directed routing technologies, it becomes immediately clear that we are facing a tussle of significant magnitude at the very core of the Internet – with users demanding choice on one side and ISPs trying to maintain fragile business relationships on the other. In particular we state the following:

***Conclusion 3:*** *The already fragile set of business relationships that underly the routing fabric of the Internet will be challenged by user demands, in the form of user-driven routing technologies. Without some means of rationalizing the economic interests of the ISPs with the desires of the users, the Internet will suffer from decreased quality and increased cost.*

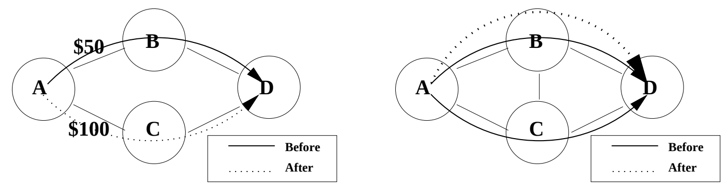


(a) Priced Service (b) Peered Networks

**Figure 1: Routing Examples**

## 5.1 Examples

To crystallize the problems and provide reference points, we present three simple examples.

*Example 1*: In Fig. 1a, AS A is a customer of both B and C to reach a set of destinations D. The price for B is \$50/Mbps but \$100 for C; thus A uses B. Now assume that a significant fraction of A's users wish to travel through C to reach D. A obviously has significant disincentive to allow this. In this situation, it is likely that A will block any form of user directed routing at all, if it has the ability to do so. If it does not have this ability, then it is clear that A will suffer.

Neither outcome is a good one.

*Example 2*: Example #1 relied on price differences to create the tension. However, many networks are peered (settlement free). Consider the similar example of Fig. 1b. Here A, B, and C are peers, with peering traffic ratios close to 1.0. Let us now assume that A's users can and wish to direct their traffic through B.

After some time A, B, and C will notice that their peering ratios are now significantly out of balance. If A is the smaller ISP it may now be forced to pay B for its future traffic – or worse it may be forced to pay *both* B and C. Knowing this, A will attempt to redistribute the portion of its traffic that is not user-directed. This solution comes with operational costs and causes traffic to be sub-optimally routed. Beyond this, a feedback look is created; the poor routing may incent users to depend further on user-directed routing, worsening the problem.

Again, the outcome is painful for A, A's users, or both.

*Example 3*: In this example, we consider traffic flowing into D in the network of Fig. 1b. Assume that D's outbound traffic is higher than its inbound (e.g., an ISP with webservers) and D is thus a customer of B and C. Assume that B and C are peers.

Now introduce overlays that choose to shift traffic from A to D going through C to instead go through B. This may level out D's traffic ratio with B, causing D to propose peering with B. To avoid this, B may try to route some of its traffic destined to D through C, defeating the purpose (or increasing the use) of the overlay. Note that this behavior, triggered by the presence of the overlay, contradicts the conclusion of Gao and Rexford.

the value of the network but also increases its cost structure.

[5] This is a critical difference from IP Multicast, another technology that posed an economic threat to ISPs.

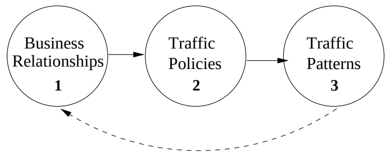


**Figure 2: A Representation of the Gao and Rexford Model with the Implicit Feedback Loop Not in Their Model**

## 5.2 Analysis

Given these examples, we now consider the issue in a more structured way, by introducing user-directed routing to the model of Gao and Rexford [1] [14]. We are interested in what happens when user-directed routing is without an additional incentive-based mechanism.

A simplification of their model, depicted in Fig. 2 is:

1. Commercial relationships are formed
2. Based on (1), traffic policies are designed
3. Traffic flows in accordance with (2)

Gao and Rexford observe that in such a model it is a reasonable traffic policy to route traffic through customers before peers or providers. They also show that this game, the Stable Paths Problem, has a solution, as discussed in the Introduction.

We first consider the extreme model of fully user-directed routing. This is not supported in Gao and Rexford's model and in fact severs the link between (2) and (3). It is unlikely that the users will coincidentally choose the paths the ISP would have chosen, leading to the ISP problems and responses described in the examples.

Given the intuition of stressing the link between (2) and (3), we now examine a weaker version of user-directed routing: overlays running over BGP. Consider the stage game of a specific moment in time. Here, the business relationships are fixed and the overlay traffic patterns, while abnormal, are fixed over BGP-defined routes. Consequently, the model and its conclusions still hold. We have:

***Conclusion 4:*** *The stage game of the Stable Paths Problem is stable even with overlay routing*

This analysis does *not* hold when the problem is played over time – that is, as a repeated game. First, the overlay removes the direct causality between steps (2) and (3). Second, step (1) in the model assumes *exogenous* business relationships. In practice, however, these relationships are based on the traffic levels themselves. Consequently, there is a feedback loop, depicted in Fig. 2, that must be considered. If an overlay controls sufficient traffic, it can (as seen in the examples) cause an AS to regret its past decisions made – even the decision of customers over peers (precisely what is demonstrated in Example #3). Thus:

***Conclusion 5:*** *The repeated game of the Stable Paths Problem where contracts are based on traffic patterns is not stable with user-directed routing.*

The contrast in our conclusions arises due to differences between our assumptions and those of Gao and Rexford. In particular we assume a) that user-directed routing enables traffic patterns to change with significantly greater magnitude and fluidity, and b) the existence of a feedback link between (3) and (1). In reality, this link exists today, but operates at the (currently longer) timescale of traffic mobility.

The effect of our more complete assumptions is that a model with user-directed routing cannot be shown to necessarily satisfy the stability property identified by Gao and Rexford over time. This is equivalent, in a more formal model, to showing that the game is not solvable in dominant strategies. This is unfortunate, because, as discussed earlier, the stability property leads to several characteristics valuable for Internet routing.

Although this strong version of stability does not exist in our model, it is not the case that the Internet must therefore plunge into massive instability as user directing routing becomes widespread. We suggest that with appropriate mechanisms and models for these conditions, equally strong notions of stability can be defined for overlay-augmented systems. We discuss this further in Section VII.

## 5.3 Potential Outcomes

Given the above structure and examples, it is not hard to imagine the set of possible outcomes over time:

- **ASes "Win"** ISPs may be successful in preventing the realization of overlays and the establishment of any other user-directed scheme. Here user-quality will suffer, these preventive practices will create cost for ISPs, a lot of research effort will have gone to naught, and the benefits will go unrealized.
- **The Users "Win"** Here (as is the case today) overlay technologies allow end-users complete flexibility in picking their routes without any economic considerations. As the prevalence of overlays increases, ISP profitability and the stability of inter-ISP relationships will be significantly degraded, in turn affecting end-users. Again, this is not desirable.
- **A New Hybrid Solution Emerges** This is the desirable outcome, but also the most nebulous. In particular, it must have the property of allowing some user choice while finding a means to appropriately compensate the ISPs for the decisions made.

We observe that unlike the undesirable solutions at either endpoint of the spectrum, "the" hybrid solution is in fact a range of possible solutions, with different solutions emerging over time and different solutions being appropriate at different points in the network. Further, we have discussed several times that the root cause of these problems is that the incentives in the system are implicit, not explicit. Based on this observation, we argue that the most effective possible path forward is to make the incentives implicit in today's model *explicit* in the routing information dissemination fabric, and allow the most appropriate hybrids to emerge as and when appropriate.

***Conclusion 6:*** *Introducing incentives, represented by prices, natively into the routing fabric will allow us to both resolve the conflict between ISPs and user-choice routing, and to address significant practical problems with BGP-based routing as it exists today.*

# 6. RESOLVING THE TUSSLE

We turn now to the characteristics of the solution called for in Conclusion #6. We first consider properties desirable for any such solution and then sketch several possible implementation strategies.

## 6.1 An Ideal Framework

We examine the properties of an ideal implementation framework. The questions below are not exhaustive nor requirements, in fact we will see that tradeoffs exist. However, they provide us with guidance and metrics.

1. *"What is the good to be priced?"* Several properties are desirable. The good should be unambiguous and easily audited, to increase the likelihood of successful transactions and minimize overhead. Its definition should be directly relevant to both the end-user's utility function and the ASes cost, so that one or both parties can easily reason about it. (This tradeoff will be significant later.)

2. *"How and when should the information regarding the good and the prices be conveyed?"* Great variety is possible, from once a year through written contracts to every second in a routing message. We identify several guiding principles. First, the time frame should be sufficiently short that the underlying economics and incentives are stable within a given period – when an entity publishes its information, it should not worry that significant changes will cause it to regret its decisions *ex-post*. Secondly, the time frame should be sufficiently large that the system can achieve an appropriate level of stability. Third, for reasons of fate-sharing, consistency, and efficiency it is advantageous if market information about the goods is exchanged in the same framework used to convey other (e.g., technical) information.

3. *"Who are the users?"*. Thus far we have spoken of ASes and end-users. However, there is a continuum from individual users to large end-users (e.g., corporations or universities) to small ISPs to large ISPs. Furthermore, there are other players, such as 3rd party overlays. Given this melange of entities, it is unclear where different incentive-response mechanisms should be placed. In answering this question, we offer the following properties. First, the user must have sufficient information to make a decision. Second, the supplier must have sufficient ability to implement the decision. Third, the benefit to the user of being able to make the decision must outweigh the cost and/or uncertainty of having to make that decision. Fourth, the benefit to the supplier of enabling the decision must outweigh the cost and/or uncertainty of facilitating such a decision. Based on these properties, it is clear that – different hybrid models, in the language above – may exist in different locations in the system.

## 6.2 Application to Routing Frameworks

We turn to the question of implementing these goals in a routing framework. We look at three frameworks – BGP as it exists today, a model of complete user/source-based routing, and a model in which overlay networks interact with a route discovery protocol. In each we see that we must provide slightly different answers to the above questions.

There are two common themes in our instantiations of our principles. First, we use the notion of *routes* to represent the good in the system. Routes are easily auditable, directly implementable by the AS, and clearly tied to the AS'es cost structure. An alternative would be a good tied to a metric of quality, along the lines of DiffServ [15]. There are several reason why we select routes as our good. First, it does not require the explicit definition of quality classes. Second, it is both easily auditable and easily implemented by the ISP, obviating the need for complicated Service Level Agreements (SLAs). Third, it maps directly to the ISP's cost structure.

The second commonality in these instantiations is our answer to when and how. Cost structures are constantly changing as contracts are renewed and the underlying topology changes. An exchange that is on the timescale of more than days will likely not facilitate stability. Therefore, all of our proposals have granularity on the order of at most hours.

These proposals are *not* intended to be complete solutions, particularly since they do not incorporate many important subtleties. Instead these solutions serve as proofs-of-concepts and motivate some important research questions.

### 6.2.1 A Next Step on BGP

The first framework we consider is that of BGP modified to incorporate the prices into the route advertisements. More specifically:

- Every inter-domain BGP route advertisement will carry with it an associated price representing a per-GB transfered price. As a matter of practice, these prices can be changed only once ever hour.
- The business relationships continue to be pair-wise between ASes, with charges now based (in part) on these prices.
- ASes can incorporate this information into their routing decisions and perhaps route solely on these prices.
- ASes can provide this information to any overlay system operating in its network.

We note several relevant implementation details. As discussed, we use routes as our logical good, but for compatibility with BGP, we represent routes as destinations since BGP enforces a one-to-one mapping at the inter-domain level. Leveraging BGP also facilities the exchange of information without a new protocol. This is true both inter-ISP and between ISPs and major customers, today many ISPs maintain BGP sessions with commercial overlay providers. We also note that the per GB pricing is consistent with average-usage billing, a popular billing methodology today. Together, we conclude that such a scheme could be easily implemented by ISPs and could be deployed incrementally at the granularity of routes.[6]

### 6.2.2 Source-Based Routing

Next, we assume a framework where a source routing protocol is used to decide among different routes. To apply an incentive scheme we propose the following:

- Every AS associates a price with each border ingress/egress pair.
- These prices can be updated on the order of minutes.
- All information on routes and associated prices are distributed throughout the network within the routing protocol.
- Each user selects the path that maximizes her utility, given the observed quality and price of each route.
- The ASes along the path obey the requested path.

[6] One subtle downside is that this could potentially cause deaggregation.

- ISPs are compensated for the use of their routes.[7]

The primary difference between the source-based and BGP-based schemes is the significant increase in information and flexibility provided to the user by the assumed protocol. Since we are not worried about the convergence of some underlying routing protocol, we can increase the frequency of price updates. Despite these differences, we again see that given the particular routing framework, we are able to infuse an incentive framework with minimal alterations.

### 6.2.3 An Overlay Controlled Environment

Lastly, we consider an in-between and perhaps more likely reality, where overlay networks and ISPs work together to provide efficiency and scalability. The first part of the routing mechanism is a system in which path existence and pricing information is propagated through the network at some relatively low frequency. Like BGP, paths are built up AS by AS. Unlike BGP, multiple paths can be advertised and changes in link status do not necessitate a corresponding advertisement. This is because the second part of the mechanism is an overlay-based technology that chooses the optimal route based on the set of paths available, their relative financial cost, and their relative quality.

Here we see the following:

- Every AS advertisement also includes a price.
- This route information is updated on the order of hours.
- Overlays, based on the information at hand and users' desires make the appropriate decisions.
- Overlays may exist as separate entities (e.g., a Content Delivery Network (CDN)) and have flexible relationships with end-users

## 7. RESULTS

Based on the above applications of the incentive scheme we now analyze the question *"Is it worth it?"* In particular, is the main problem of the tussle and its impact on stability really resolved? Second, have we addressed the problem in a way that is not excessively burdensome to the players or the market?

First we examine our impact on the tussle itself. Our schemes have made the incentives of the ASes, currently implicit, explicit to each other and to end-users. Furthermore, they have transformed peering relationships from implicit to explicit relationships. Since we showed that the implicit nature of the relationships was the root problems in all three examples of Section V, it is clear that our schemes resolve the problems. Beyond this, it is possible to argue, under similar assumptions to Gao and Rexford, that we now have *ex-post* stability even in the repeated game. By re-introducing stability and resolving the conflicts, we create a framework in which ASes are willing to support user-directed routing.

Now we examine the costs at which these benefits have come. One potential downside is that we could have introduced significant complexity into the system through new or modified protocols. The use of routes as the good and the leveraging of protocols that already deal with routes allows us to suggest, in Section VII, that we did not do this. Another source of cost is that the ISPs must now track usage with finer granularity. However, this can be implemented solely at the ISP border, is becoming more supported in routers [16], and can be limited to those routes where the added monitoring is worthwhile. We note that an ISP must already today monitor the traffic of its peers, to ensure that it is making the right peering decisions. Finally, we note because the business relationships are now in-band, the security, robustness, and auditability of the routing system becomes even more important. However, these points are already of great importance today. Nonetheless, we discuss some of these points in the next section.

Finally, in response to the problems of Section III, we argue that this scheme makes routing in a BGP-framework *simpler and more efficient*, even in the absence of user-directed routing. While ASes may (and likely will) continue to implement a lower or lowest-cost routing policy, the clarity of incentives will prevent perverse routing pathologies designed to maintain odd business relationships. Furthermore, the complexity of reducing cost through inter-domain routing is significantly reduced, which in turn can decrease the cost structure of ISPs.

## 8. REALIZING THE MECHANISM

Having seen the potential benefits and practicality of such a mechanism, we can examine several other key and open interesting research questions.

- *Who is the user?* In Section V-C we pointed out that there exist a range of potential hybrid mechanisms. In Section VI-B, we presented a spectrum of answers and we believe that a continuum of implementations is not only optimal but also presents a plausible adoption path. In particular, it is unclear that every end-user will want to be making these decisions. Thus, we suspect that the "end-user" in our models will primarily be the access ISPs and/or the overlays. These entities in turn can have relationships with end users where the tradeoffs are more manageable or well understood. Note that this issue is intimately related to the question of how to make the system scale – an important question for any user-directed routing system.
- *Who absorbs the uncertainty?* Building out or maintaining capacity is not cheap, and there already exists a mini-tussle between players seeking longer-term contracts and players seeking more flexibility. User-directed routing brings this more into focus. Uncertainty over traffic volumes will exist in any system – but who should absorb the uncertainty? Should we look to 3rd parties such as CDNs? One compromise solution would be to employ user-directed routing on select paths or for a fraction of one's traffic (e.g., the important flows). However, even this simple approach will only be realized if the users are properly incented.
- *What will the steady-state dynamics be?* One issue underlying many of these questions is how dynamic traffic patterns will be with user-choice routing. We see this as three questions. First, "How much traffic will want to move?" This has been addressed in part by studies such as [4]. A second question is "How much traffic will want to move *again*?" That is, how dynamic will traffic be

[7]The means by which the ISPs are compensated is critical to the success of the implementation. Two possibilities are that the user pay each ISP along the path or that payments are accumulated pair-wise.

in a user-directed routing regime. Perhaps most important, "For what fraction of traffic that desires to move will the benefits of the move outweigh the costs of doing so?" The commercial success of CDNs suggest that resulting number is still significant, but (by definition) the two latter questions each reduce this quantity. Further research here would be fascinating.

- *What form of pricing model(s) and business relationships are appropriate?* In these systems, it is desirable that the pricing scheme be a) simple and b) composable (or more formally, distributive). Therefore, in our proposals we used a volume transfered scheme, which satisfies both properties. Another popular model, 95/5, satisfies the first cleanly but less completely satisfies the second. Since these models are not perfect even without user-directed routing, envisioning other possible models is interesting.

  Given a means of calculating cost, we see that customer-provider relationships are simply implemented. But what of peers? We can treat "peer" relationships as two customer-provider relationships, with the assumption that good peering relationships will lead to a net payment sufficiently near zero. Also, we need a clear way of determining payments when these "peers" need to pay one another, such as a simple sender-pays rubric. This is in contrast to today where we have *both* sender-pays and receiver-pays aspects.

- *Incentive Compatibility* We have simplified the logic for forwarding, but have not addressed price setting. In [17], Feigenbaum *et al* address this question and present a strategyproof mechanism based on BGP. However, there are several strong assumptions in their model. Can we relax any of these while minimizing the strategizing of the players?

- *What of privacy?* Today networks generally try to keep the details of their business relationships private. This is greatly reduced in an explicit incentive model. An interesting question is *"How much privacy is really lost in going from today's implicit model to an explicit model?"*. For example, Subramanian *et al* show how many business relationships can be inferred from public BGP feeds [18]. Furthermore, NDAs are known not to be perfect. Thus the question can be addressed from a theoretical, a modeling, and/or a practical perspective.

- *What are appropriate models?* Gao and Rexford's model, while quite simple, was very valuable to this paper. It provided both a framework and the important property of stability. Further, it provided a structured means of explicating the impact of user-directed routing.

  We however also saw that this model was not perfect, raising the question "*What types of models will be most useful in reasoning about incentive-based routing systems?*" There are a massive number of parameters to consider. We offer a few points. First, we feel strongly that instead of invoking the nebulous notion of (heterogeneous) policy, *money – a universal motivating factor – should play a central role.* Second, heterogeneity should be pushed out toward the edge, where users' preferences, particularly among applications, vary widely. Furthermore, we believe that there are gains to be made from bridging the gap between protocol specific models and idealized models of routing, such as the ones used in the Price of Anarchy analysis.[19] [20]

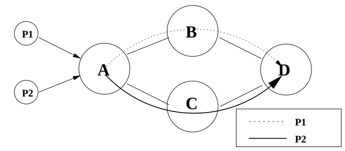


**Figure 3: Price of Tyranny with the Optimal (Non-Confluent) Paths Depicted**

# 9. WHAT IS THE PRICE OF TYRANNY?

In [21] Papadimitriou *et al* ask the question "What is the cost of the architecture of the Internet?" They and others have then gone on to examine "Price of Anarchy," the ratio between the aggregate utility in a network run by a socially optimal planner and one in which selfish users can make their own decisions.

We examine another cost of the Internet's architecture within the framework of this paper and based on the second modeling property proposed above. Despite the fact that user's preferences are different (or perhaps different for different types of traffic), traffic on the Internet is confluent. Thus, even if the ISPs were solely interested in maximizing the users' aggregate utility, the result would be worse than in the case of a complete user-choice model. Therefore, we propose, model, and analyze *The Price of Tyranny.*

## 9.1 The Problem

In this model each link has its own quality and cost metrics. To this network we add a set of users, each with preferences over (quality, cost) pairs. More formally:

- A non-directed, connected graph $G = (V, E)$ with no self-edges
- A penalty $p_{(i,j)}$ and a cost $c_{(i,j)}$ $\forall (i, j) \in E$
- A user-base consisting of a set of users $N$ and connecting-edges $L$ such that $\forall i \in N$ $\exists (i, j) \in L$ where $j \in V$
- A route, $r$ through $V$, is a set of edges $r \subseteq \{E \cup L\}$
- The cost of a route $c(r) = \sum_{(i,j)\in r} c_{(i,j)}$
- For each player ($i \in N$), a quality function over routes $q_i(r) = \hat{q}_i(e_1, e_2, ...., e_d)$, $e_x \in r$ which is non-decreasing in the sum of the penalties.
- A traffic incidence matrix $T_{(i,j)} i \in N, j \in V$ .
- For each user a utility function $u_i(r) = \hat{u_i}(q(r), c(r))$
- A map $M$ is a set routes $M(i, j)$ such that $\forall i \in N, j \in V$ $\exists r_{(i,j)} = \{(i, y_1), ..., (x_d, j)\} \in M$.
- The mechanism is individually rational. That is, $i$ sends to $j$ iff $u_i(q(r), c(r)) \geq 0$, $r = M_{(i,j)}$.
- The utility of a map $U(M) = \sum_{r_{(i,j)}\in M} T_{(i,j)} \cdot u_i(q(r), c(r))$

We also define $C$ to be the set of confluent maps – for a given destination, every node has exactly one next hop, as is the case on the Internet.

Let $M_{UC}$ be the map where each user chooses its routes. Let $M_{BC} = \max_{M\in C} U(M)$, that is the best confluent map.

We therefore define the Price of Tyranny, $\tau = \frac{U(M_{UC})}{U(M_{BC})}$.

| | P1 | P2 |
|---|---|---|
| Top | 10 | 5 |
| Bottom | 0 | 10 |

(a) Scenario #1

| | P1 | P2 |
|---|---|---|
| Top | 10 | 0 |
| Bottom | 0 | 10 |

(b) Scenario #2

**Figure 4: Game Payoffs**

## 9.2 Analysis

Larger $\tau$ implies a larger benefit for incentive-based routing but is unfortunate for the Internet today. Therefore, we seek to find upper-bounds for $\tau$ under reasonable assumptions. To gain intuition, we examine the simple graph in Fig. 3 with corresponding payoffs in Table 4. In Scenario #1, an omniscient and benevolent AS chooses the top route. This has a utility of 10+5=15 yielding $\tau = \frac{10+10}{10+5} = 4/3$. However, in Scenario #2 we have $\tau = \frac{10+10}{10+0} = 2$. In fact, we can show that for two users and two paths this is the worst case.

***Lemma #1*:** On the graph in Fig. 3, $\tau \leq 2$.

*Proof Sketch*: This follows from maximizing the ratio subject to the constraint that the path selected is socially optimal.

Furthermore, it can be shown that $\tau$ is *independent of the number of users*, a surprising and positive result:

***Theorem #1***: On a general graph, $\tau \leq |E|$.

*Proof Sketch*: This follows from examining sets of users, where users in the same set prefer a given route in the graph.

This result is surprising and positive as it shows that we do not lose efficiency as the Internet's user base grows. However, $\tau$ is dependent on $|E|$, or more precisely the number of unique routes in the graph. Thus, we feel that tighter upper bounds on $\tau$ are possible, especially if we assume structure on the graph and/or bounds on the differences between utility functions. Lower bounds under reasonable assumptions would also be interesting. Correlating quality and load (as in the Price of Anarchy work) would be another interesting extension.

# 10. SUMMARY

In this paper we presented and analyzed the tussle between user-directed routing and ISPs, and demonstrated how this results from the current routing mechanism. We proposed the notion of incorporating prices into the routing system; and demonstrated, through applications to various routing architectures, that this can be achieved with minimal technical steps and may instead simplify the system. We concluded by presenting several open research questions.

# 11. REFERENCES


[1] Lixin Gao and Jennifer Rexford. Stable internet routing without global coordination. In *Measurement and Modeling of Computer Systems*, 2000.

[2] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden. Tussle in cyberspace: Defining tomorrow's internet. In *ACM SIGCOMM 2002*, Pittsburgh, PA, August 2002.

[3] Renata Teixeira, Aman Shaikh, Tim Griffin, and Jennifer Rexford. Dynamics of Hot-Potato Routing in IP Networks. In *Proc. ACM SIGMETRICS*, 2004.

[4] Stefan Savage, Andy Collins, Eric Hoffman, John Snell, and Thomas E. Anderson. The end-to-end effects of internet path selection. In *SIGCOMM*, pages 289–299, 1999.

[5] Nick Feamster. Practical verification techniques for wide-area routing. *SIGCOMM Comput. Commun. Rev.*, 34(1):87–92, 2004.

[6] Andrew M. Odlyzko. Pricing and architecture of the internet: Historical perspectives from telecommunications and transportation. http://www.dtc.umn.edu/ odlyzko/doc/-pricing.architecture.pdf.

[7] Tobias Oetiker. Mrtg: The multi router traffic grapher. In *Proceedings of the 12th Conference on Systems Administration*, pages 141–148. USENIX Association, 1998.

[8] Akamai technologies homepage. http://www.akmai.com/.

[9] Dave Andersen, Hari Balakrishnan, Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, 2001.

[10] David D. Clark. Policy routing in Internet protocols. RFC 1102, Internet Engineering Task Force, May 1989. ftp://ftp.ietf.org/rfc/rfc1102.txt.

[11] Xiaowei Yang. Nira: a new Internet routing architecture. *SIGCOMM Comput. Commun. Rev.*, 33(4):301–312, 2003.

[12] Alex C. Snoeren and Barath Raghavan. Decoupling policy from mechanism in internet routing. *SIGCOMM Comput. Commun. Rev.*, 34(1):81–86, 2004.

[13] Akamai glitch causes web traffic disruption. http://biz.yahoo.com/rc/040524/-tech_akamai_outage_2.html.

[14] Lixin Gao, Timothy Griffin, and Jennifer Rexford. Inherently safe backup routing with BGP. In *INFOCOM*, pages 547–556, 2001.

[15] An architecture for differentiated services. http://www.ietf.org/rfc/rfc2475.txt.

[16] Interpret the destination class usage mib. http://www.juniper.net/techpubs/software/junos/-junos56/swconfig56-net-mgmt/html/dcu-mib.html.

[17] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, and Scott Shenker. A BGP-based mechanism for lowest-cost routing. In *Proceedings of the 2002 ACM Symposium on Principles of Distributed Computing*, pages 173–182, New York, NY, July 2002.

[18] Lakshminarayanan Subramanian, Sharad Agarwal, Jennifer Rexford, and Randy H. Katz. Characterizing the internet hierarchy from multiple vantage points, 2001.

[19] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[20] Ramesh Johari and John Tsitsiklis. Routing and peering in a competitive internet. citeseer.ist.psu.edu/572286.html.

[21] Christos H. Papadimitriou. Algorithms, games, and the Internet. *Lecture Notes in Computer Science*, 2076, 2001.