Analysis of Inter-Domain Collaborative Routing: Provider Competition for Clients

Martin O. Nicholes, Chen-Nee Chuah, Shyhtsun Felix Wu, and Biswanath Mukherjee

Abstract: **Any server offering a routing service in the Internet would naturally be in competition for clients, and clients may need to utilize service from a specific server in order to achieve a desired result. We study the various properties of this competition, such as the fraction of route requests handled by a routing service provider and the fraction of total revenue obtained. As the routing service providers (i.e., servers or routers in this context) compete, they may alter behavior in order to optimize one of the above properties. For example, a service provider may lower the price charged for its service, in order to increase the number of clients served. Our models are based on servers offering a routing service to clients within representative network topologies based on actual Internet sub-graphs. These models provide a framework for evaluating competition in the Internet.**

We monitor key aspects of the service, as several variables are introduced into the models. The first variable is the fraction of client requests that will pay more for a better quality route. The remaining requests are normal client requests that are satisfied by the most economical route. The second variable is the fraction of servers who choose to lower service prices in order to maximize the number of client requests served. As this fraction increases, it is more likely that a server will lower the price. Finally, there are some resource constraints applied to the model, to increase the difficulty in providing a routing solution, i.e., to simulate a realistic scenario.

We seek to understand the effect on the overall network, as service providers compete. In simple cases, we show that this competition could have a negative impact on the overall efficiency of a service. We show that the routing variety present in the larger models is unable to mask this tendency and the routing service performance is decreased due to competition.

Index Terms: **Collaborative routing, competition, inter-domain routing, pricing, routing.**

I. INTRODUCTION

Servers providing a routing service in the Internet will naturally be in competition for clients, and clients may need to utilize a specific server to meet some requirement. Success for one client could mean failure for another client, if a necessary server or link runs out of capacity. Our models are based on servers offering a routing service, called inter-domain collaborative routing (IDCR) [1] to clients within representative network topologies based on actual Internet sub-graphs. An IDCR server provides a routing solution to clients seeking to avoid a single autonomous system (AS) in the network. An AS is equivalent to an entity in the Internet, such as a corporate domain. (IDCR will be described briefly in Section IV.) These models provide a framework for evaluating competition in the Internet.

We measure various aspects of this competition, such as the fraction of routes provided by an IDCR service provider and the fraction of revenue obtained. The fraction of routes provided is a measure of the 'market share' of the IDCR service provider. The term IDCR server is used to denote an IDCR service provider.

We monitor key aspects of the service, as several variables are introduced into the models. The first variable is the fraction of premium routes. Premium routes simulate a client route where incremental hops of the routing solution are minimized, within the limits of a budget. Normal routes simply optimize the price of the route within the cost budget. The second variable is the fraction of servers who choose to lower, or discount, the prices for offered routes by a certain percentage. As this fraction increases, it is more likely that a server will discount the price.

As clients compete for limited server and network resources, the effect on the probability of routing success, $P(s)$, is important to measure. We show using simple models that the $P(s)$ can be negatively impacted by competition. However, in larger models, this effect could be masked by the large amount of route diversity available to IDCR servers and their clients. When clients must choose routing solutions within budgetary constraints, our study shows how $P(s)$ can increase as more IDCR service providers lower prices, allowing more clients to purchase services. However, as more servers discount services, some clients may lose routing services.

Finally, there are some resource constraints applied to the model, in order to simulate a realistic scenario. The resource constraints are the number of connections an IDCR server can maintain, and the capacity of the links in the network.

We seek to understand the effect on the overall network, as IDCR servers (i.e., servers or routers in this context) compete. Various metrics are measured in the models for each IDCR server, such as: The fraction of routes provided and the fraction of total revenue received. In simple cases, we show that this competition could have a negative impact on the overall efficiency of the IDCR service. We show that the routing variety present in the larger models is unable to mask this tendency and the IDCR service is affected.

The rest of the paper is organized as follows: Section II covers related work, Section III reviews the models used to study the effect of competition in a network, Section IV outlines the IDCR algorithm used in the study, Section V covers the method-

Manuscript received June 25, 2010; approved for publication by Young Han Kim, Division III Editor, June 28, 2011.

The work was performed while M. O. Nicholes was a Ph.D. student at UC Davis. He is currently employed at Insyde Software.

M. O. Nicholes and C.-N. Chuah are with the Department of Electrical and Computer Engineering, University of California, Davis, CA, 95616, USA, email: {martyon, chuah}@ucdavis.edu.

S. F. Wu and B. Mukherjee are with the Department of Computer Science, University of California, Davis, CA, 95616, USA, email: {wu, mukherje}@cs.ucdavis.edu.

1229-2370/11/$10.00 © 2011 KICS

ology used to study competition among routers that provide the IDCR service, Section VI reviews and discusses the results of the study, and Section VII concludes the paper.

II. RELATED WORK

Our work on evaluation of the effects of competition on inter-domain routing falls into the area of applying game theory concepts [2] to inter-domain routing. A survey covering the use of game theory in the study of communication networks stated succinctly that game theory allows algorithms to be studied when network nodes are not fully cooperative [3].

One early work investigated properties of inter-domain routing with 'selfish users,' in the context of game theory [4]. The work showed that, in some simple networks, the Nash equilibrium can be proven to exist. The work classifies the inter-domain routing problem with selfish users as a noncooperative game. Our work attempts to measure effects of these types of games in more representative networks, discussing Nash equilibrium from the server perspective.

Another excellent work explored game theory and network pricing from a multi-class network service perspective [5]. Similar to our work, the paper proposes two pricing tiers, one for 'flat pricing,' and another for 'priority pricing.' Simulations were performed with a mixture of application types to determine a pricing scheme that would be 'adoptable,' i.e., satisfactory to end-users. Our work is performed on larger, more representative network models and focuses on client-server interaction in the area of network routing.

Another work reviewed Internet pricing using a quality of service (QoS) model [6]. The work chose to view the problem from a cooperative game perspective. In a cooperative game, there is a bargaining strategy used by the 'players.' In addition, the degradation of network services, as more services are sold, is investigated at the Internet service provider/customer interface. Our work applies the concept of strategy to IDCR routing services. We seek to demonstrate that degradation of services can occur in larger network models, as a result of server strategy choices.

In a related work, the competition of network access providers (NAPs) for network resources is evaluated, as the NAPs try to optimize their own performance [7]. The work determines conditions in this network competition, where a Nash equilibrium point (NEP) can be determined. At the NEP, no NAP can increase its performance by deviating from their current strategy.

Competition of inter-domain routing in optical networks is explored in another work, using constrained network models [8]. The goal of the end-user is to route traffic cheaply, while each domain, or player, attempts to provide routing service for the traffic. The links in the model have limited capacity, so a single domain cannot route all the traffic. This is very similar to our work, although our work attempts to evaluate the effect of competition in larger, more-realistic models, and tracks incremental hops incurred as clients use alternate routes.

Competition between ASes who are servicing a multi-homed client is investigated, as this related work reviews possible service provider techniques, including AS path manipulation and AS prefix hijacking [9]. These techniques are non-standard for

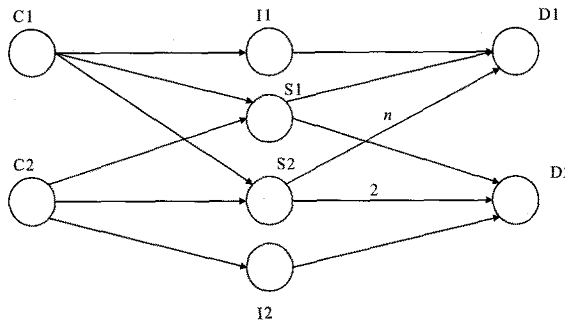

Fig. 1. Simple network for illustration purposes.

an AS to employ, but are shown to increase the AS's share of client traffic. Our work focuses at a higher level and seeks to understand the side-effects of competition in a constrained and standards compliant network.

A similar work looks at a model where nodes in the network must bid for network traffic, and also review bids from upstream nodes to determine how to route traffic [10]. The work proposes a link cost function, that causes nodes to route traffic, using multiple relays, to minimize link cost as well. Our work simply sets constraints on link capacity, and clients either accept or reject offered routing services from each relay server.

Pricing of network services within a QoS framework is a similar problem [11]. The work discusses a framework for dynamically pricing QoS transit services and proposes the iterative allocation adjustment heuristic, which helps an ISP assign bandwidth to various services for optimal revenue. In addition, the paper offers a method for a network engineer to determine the 'value' of a network link. Our work also deals with network service pricing, but addresses the problem from a competitive point of view.

III. NETWORK MODELS

A. Effect of Competition in a Simple Model

In this section, we show the effects of limited resources and competition on routing in a simple network, shown in Fig. 1. In the figure, clients have a C prefix in the label, IDCR servers have an S prefix, destinations have a D prefix, and intermediate nodes have a prefix of I. All links have a weight of 1, except for those marked. The link weight in the graph is an abstraction for the number of hops between the nodes. In other words, there are nodes internal to the links. Link S2-D2 has a weight of 2, and link S2-D1 has a weight of h, where $h > 2$.

The default route for C2 to reach D2 is C2-I2-D2 and is 2 hops. If C2 wishes to avoid I2 (perhaps I2 is not trusted), the alternate route C2-S1-D2 could be used, which is also 2 hops. Another possibility for C2 is to use the route C2-S2-D2, but that route is 3 hops. If C2 chooses to use the alternate route through S1 due to the lower hop count, the connection takes up resources in S1. The default route for C1 is C1-I1-D1 with an alternate route of C1-S2-D1, which costs $h + 1$ hops. If C2 has already made its connection, then C1 must use S2, since S1's resources are assumed to be fully consumed by C2's route. So, if $h > 2$, then C2's choice of server S1 causes the overall system cost to

Table 1. Route discounting game using the simple network.

		Server 2	
		Discount	No discount
Server 1	Discount	1, 1	3/2, 1
	No discount	1, 3/2	2, 2
Payoffs to: (S1,S2)			

be non-optimal (alternate route costs higher by $h-2$). If C2 reserves the alternate route before C1, then client C2 is said to have 'first mover advantage'. The optimal move of the first client places the second client at a disadvantage, and places the network in a sub-optimal state. In this scenario, no rerouting is allowed.

In order to demonstrate the impact of competition on the $P(s)$ of the network, imagine that link S1 - D1 is removed from the network. Further, assume that, in this simple model, S1 can handle only a single connection. In this situation, if C2 consumes the resources of S1, then C1 has no possible alternate route to reach D1. Since C2 optimized hop count when selecting S1, this selfish choice has left C1 without a viable route. This cuts the $P(s)$ for the overall network in half.

The simple network can also be used to demonstrate application of game theory to networking models. Assume there are two players in the game, S1 and S2. Also, assume that the two players (servers), S1 and S2, must choose from the following strategies: Discount route prices, or not to discount. Further, let the payoff to the servers be the charge received for handling a route. The normal price for handling a route is 1 unit, while the discounted charge is $\frac{1}{2}$. Table 1 shows the payoff received for each set of strategies employed by the servers. For example, the cell containing (1,1) denotes a payoff of 1 unit for S1, and 1 unit for S2, when both servers employed the discount strategy. The calculations are based on the following conditions: 1) S1 and S2 have capacity for three routes, 2) four routes are needed to fully mesh C1 and C2 with D1 and D2, and 3) the client will select the cheapest route or in case of equal cost, the shortest route. Using the Gambit game analyzer [12], the Nash equilibria are found to be (discount, discount) and (no discount, no discount). If one server thinks the other server will not discount, then the best payoff is obtained by matching the other server's strategy. The same is true if one server thinks the other will discount. This game theory concept will be revisited in the discussion of the results in Section VI.

This simple example demonstrates the routing problems that can be introduced due to resource constraints and competition for network services.

B. Main Models

B.1 Overview

In order to evaluate competition among routing service providers, several models were created using the scalable simulation framework network models (SSFNET) simulator [13]. There are three models utilized in the study. They all approximate the power-law graphs present in the Internet. The first model, containing 32 nodes, is shown in Fig. 2. Tier-1 routers are enclosed by the dashed oval in the center of the figure. This model is leveraged from a posted SSFNET model. During the model runs, clients of the routing service are limited to nodal degree 1. Since all the clients are of degree 1, they do not have access to any alternate paths via normal routing updates, which are received via border gateway protocol (BGP). BGP is the protocol routers use to exchange route information in inter-domain routing. Only 5 routing service providers are utilized in the 32-node model runs.

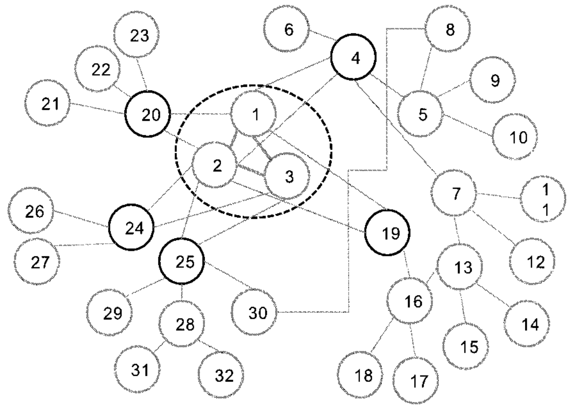

Fig. 2. Small 32-node model.

The second and third models are larger models that are used to demonstrate scaling of the observed effects and are created using a tool posted on the SSFNET web site [14]. The second model contains 110 nodes. The third model contains 208 nodes. The last two models are both based on small subgraphs contained within the actual Internet. These models do not have any nodes of degree 1, so the client nodes were allowed to have a maximum degree of 2. In addition, since clients in these models are dual-homed, i.e., connected to more than one network connection, the client has access to multiple paths via BGP updates. However, for the model runs, clients were not permitted to utilize alternate paths available in their local routing database. Instead, clients had to rely on the routing service providers for an alternative route.

Our models all contain servers offering a routing service, called IDCR [1], to clients within network topologies. An IDCR server provides a routing solution to clients seeking to avoid a single AS in the network. This model provides a framework for competition in the Internet. The model is also effective at demonstrating the effects of competition.

B.2 Measurements

We measure various aspects of this competition, such as the fraction of routes provided by an IDCR service provider and the fraction of revenue obtained. The fraction of routes provided is a measure of the 'market share' of the IDCR service provider. The term IDCR server is used to denote an IDCR service provider. The fraction of revenue obtained is important, as this determines if the market share is returning a comparable fraction of the revenue. As the service providers compete, they may alter behavior in order to optimize one of these aspects. For example, an IDCR server may lower the price charged for a route, in order to increase the number of clients served.

As clients compete for limited server and network resources, the effect on the probability of routing success, $P(s)$, is important to measure.

B.3 Variables

Several variables are introduced into the models. The first variable is the fraction of premium routes. Premium routes simulate a client route where incremental hops of the routing solution are minimized, within the limits of a budget. This type of route might be used for traffic that is sensitive to round-trip time. Normal routes simply optimize the price of the route within the cost budget. Why would a client settle for a route that is cheaper, but has more incremental hops? The traffic may be from a backup data stream, which has relaxed constraints on round-trip time.

The second variable is the fraction of servers who choose to lower, or discount, the prices for offered routes by a certain percentage. As this fraction increases, it is more likely that a server will discount the price. As the model runs, a server must decide whether or not to discount, based on comparison of a performance metric with a threshold setting. The server may have a tendency to discount, but then makes a separate runtime decision whether or not to discount each route.

B.4 Constraints

Finally, there are some resource constraints applied to the models, in order to simulate a more realistic scenario. The resource constraints are the number of connections an IDCR server can maintain, and the capacity of the links in the network. Fixing the number of connections a server can maintain is the main method to force competition among IDCR servers. Each server can only service a fraction of the requests for routing solutions. By limiting the capacity of links in the model, an IDCR server with low nodal degree may find link capacity constraints determining how many clients can be served. These constraints, even though they make our models more realistic, relate to only a portion of the constraints required to duplicate the actual network. These simple constraints are able to force routing solutions to use more links in the network and approach the type of load balancing that occurs in a real network, which was critical for our study.

IV. INTER-DOMAIN COLLABORATIVE ROUTING (IDCR) ALGORITHM

This section outlines IDCR [1], and then presents the details of the IDCR algorithm that were used in this study. IDCR is used as a sample multi-path protocol for evaluating competition between servers and clients. Another multi-path protocol could be substituted.

A. *Inter-Domain Collaborative Routing (IDCR) Algorithm Overview*

The IDCR algorithm allows IDCR servers to collaborate with client routers, if necessary, to find a route meeting some requirement, such as avoiding a particular AS. Fig. 3 clarifies the entities used in the algorithm. The IDCR algorithm performs the following actions at each client, or AS_S.

1) Determine the route, P_m, to an IP prefix, IPd, using the BGP default route,

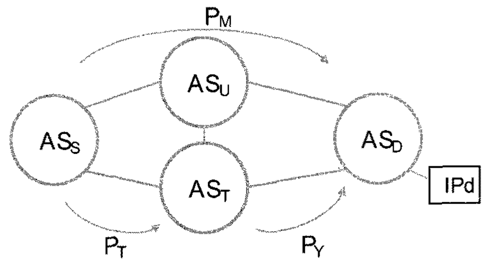

Fig. 3. Important concepts of IDCR.

2) Determine if the untrusted AS, AS_U, is on the route,
3) Determine if AS_U is on the BGP default route to the collaborating router, AS_T, and
4) Collaborate with AS_T to create a route using a broadcast request and response mechanism. The final route consists of the concatenation of the sub-path P_T from AS_S to AS_T and the sub-path P_Y from AS_T to the destination, AS_D.

B. *IDCR Modifications*

The IDCR protocol was extended slightly in order to gather the data for this work. IDCR server code was changed so that a cost would be calculated and attached to an offered route. The cost calculation is based on route quality, measured in incremental hops, as well as on the resource loading of the server. A more-heavily-loaded server will increase the route cost. In order to determine the best response from a server, the client would not only review the incremental hops of the offered route, but would weigh the cost of the route to be paid to the server.

In addition, the concept of a link load monitor was introduced. A particular IDCR load monitor was selected a priori to monitor the number of trusted connections utilizing each link. Before an IDCR server would offer a route, the server would verify that the links could handle the load, using the load monitor. After a route was accepted by a client, the server would notify the load monitor. The load monitor would then update the load on the links in the network.

V. METHODOLOGY

A. *Variables and Measurements*

A handful of routers are pre-selected to provide the IDCR service, before the models are run. In addition, the values of other controlled variables are determined before the set of runs, including: 1) Link capacity, 2) maximum number of connections per IDCR server, and 3) IDCR server discount percentage. The set of runs is then performed, varying the values of the following independent variables: 1) The fraction of servers who may discount route cost, 2) the fraction of premium routes handled by clients, and 3) the budget available for clients to purchase routing services.

Clients determine an AS_U for a route as described in Section V-B. The client, AS_S, then collaborates with the IDCR servers in the model to find a route around the single chosen AS_U.

The calculation of the cost of the route offered by the server is influenced by the fraction of servers discounting. During initialization, in order to simulate a practical setting, each IDCR server randomly determines if it has a tendency to lower route prices. This is called a *strategy* in game theory terms. When the IDCR server determines that its acceptance rate is lower than a pre-determined target value, the server with this discounting tendency will start to discount. The acceptance rate is the rate at which clients are accepting the offered routes from the server. This rate is monitored for a period, and if the rate is lower than the target, the server begins to discount route pricing.

The calculation of the cost for a route is also affected by the resource loading on the server. The intent of this relationship between the cost and the resource loading was an attempt to load balance clients across servers. As a server becomes more heavily loaded, the cost to collaborate with the server increases. The cost increase provides some incentive for clients to seek other servers.

Each run of the model provides a different level of service to client routers, measured by the probability of routing success. In addition, the side effect of incremental cost in hops to use an alternate route is tracked. Since the goal of this work is to uncover the effects of competition and collaboration, statistics are maintained on each IDCR server and client.

B. Model Runs

Each model run under SSFNET takes place in three phases. In phase one, the configuration of the IDCR service is performed, setting all the controlled variables, including the identities of the IDCR servers. This information is passed into the model run via a master SSFNET configuration file. Next, the independent variable set is chosen and passed into the run-specific SSFNET configuration file. These are off-line processes. Then, the actual SSFNET simulation is performed.

The routers providing the IDCR service advertise the service in phase two of the model run. Client routers create the set of service providers, $\mathbb{F}$, using the advertisements. The IDCR protocol appends the advertisement in the form of a route attribute added to BGP route updates. Each BGP route update includes route information, such as an Internet protocol prefix (IP prefix), and the origin AS. The AS which services an IP prefix is the origin AS for that IP prefix.

Finally, in phase three, client routers perform an iterative routing test to gather routing success statistics. The routing test requires the router to determine default routes to every destination and then attempt to route around each intermediate AS in each route. The test performs the following detailed steps at each client.

1) Select the k IP prefix/origin AS (OAS) pairs AS_{D1}/IP_{D1} to AS_{Dk}/IP_{Dk} (where k represents the number of destinations, and is typically 10-100),
2) Determine the default route to AS_{Dl}, where l represents a particular destination in the range of 1-k,
3) Determine the list of intermediate ASes, AS_{U1} to AS_{Ux}, where x is the number of intermediate ASes that could be assumed to be untrusted, and
4) Attempt to collaboratively create a route to AS_{Dl} while avoiding AS_{Uz}, where z represents a particular intermediate (and untrusted) AS in the range of 1-x.

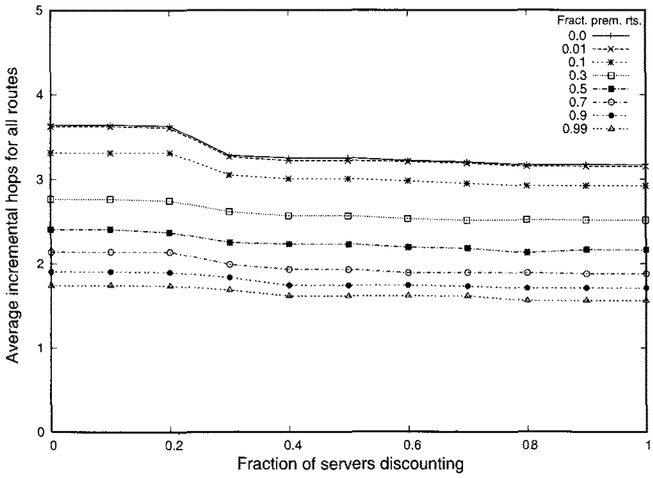

(a)

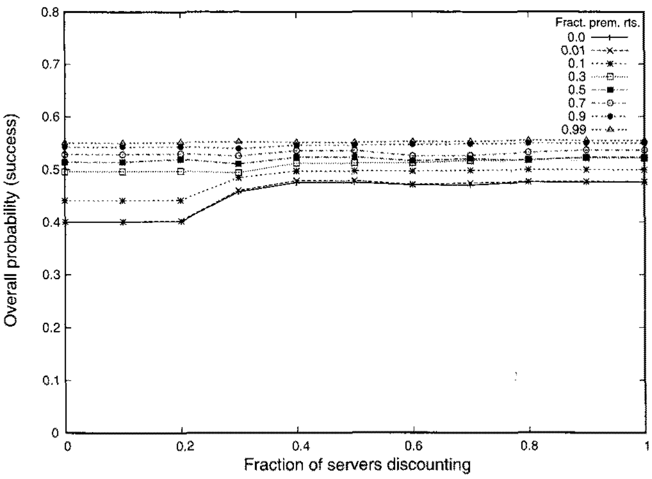

(b)

Fig. 4. Overall results with budget (3 6): (a) Overall average incremental hops and (b) overall P(s).

For example, if the route from AS17451 to AS6939 is "17451 42 6939," then AS42 would become AS_U. An attempt would be made to create a route from AS17451, AS_S, to AS6939, AS_D, that does not transit AS42, AS_U. More than one simultaneous untrusted AS is not handled by IDCR.

The model runs were executed using a version of the IDCR algorithm. We believe this scheme would be useful for evaluating the effects of competition on other inter-domain routing schemes, such as multi-path interdomain routing (MIRO) [15].

VI. RESULTS AND DISCUSSION

In this section, we show the positive effects of competition, including: an increase in $P(s)$, and a decrease in average incremental hops of a routing solution. In addition, we show that some clients may experience a loss of routing solution, due to pressure on network resources.

A. Top-Level Metrics

Figs. 4(a) and 4(b) show the most critical measurements for the overall IDCR service, namely the average incremental hops, and the overall $P(s)$. All curves are from the 110-node model.

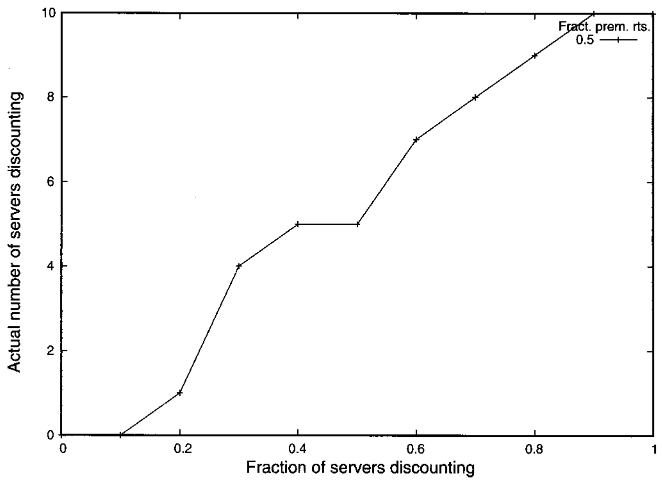

Fig. 5. Number of servers with discounting tendency.

These curves were created using data obtained with a budget value pair of (3 6), meaning that average routes could spend up to 3 units, while premium routes could spend up to 6 units. The fraction of premium routes, ranging from 0.0 to 0.99, is used as the label for each curve. For example, the curve labeled 0.5 shows results where the fraction of premium routes was 0.5.

We found that, with the spending budget of (3 6), most routers could at some fraction of premium routes purchase a routing solution within budget. In these curves, there is a general transition point that occurs at fraction of servers discounting = 0.2. At this point, a number of IDCR servers start discounting route pricing, leading to an increase in $P(s)$ and a decrease in the average incremental hops for a trusted route. This increase in $P(s)$, shown clearly on the line with x markers in Fig. 4(b), is due to the ability of more clients to afford routing solutions.

Although it appears that many of these curves have a transition in the range of 0.2 to 0.3 fraction servers discounting, we found that, by changing the random number generator seed, the transition point moved slightly, but the curves were very similar to those presented in this work.

In order to more clearly see the actual number of servers that have the discounting tendency during the model runs, Fig. 5 shows a graph of this data. For example, at fraction of servers discounting = 0.4, there are 5 servers with the tendency to discount.

B. Client Acceptance and Routes Provided

Let us examine the client acceptance curves for AS 1 in Fig. 6(a). When there is a low fraction of premium routes, i.e., most clients are simply looking for the cheapest route (lines marked 0.01 and 0.1) within the fraction of servers discounting interval of 0.1 to 0.2, AS 1 is not discounting routes, so the accept rate is low. However, when the fraction of servers discounting increases to 0.3, several more ASes began to discount, including AS 1. When AS 1, a tier-1 AS, offers a lower price, its client acceptance rate increases dramatically (0.13 to 0.22 for fraction of premium routes = 0.01). However, the fraction of routes provided by AS 1 does not respond similarly as shown in Fig. 6(b). There is only a weak increase in the fraction of routes provided at fraction of servers discounting = 0.3. There

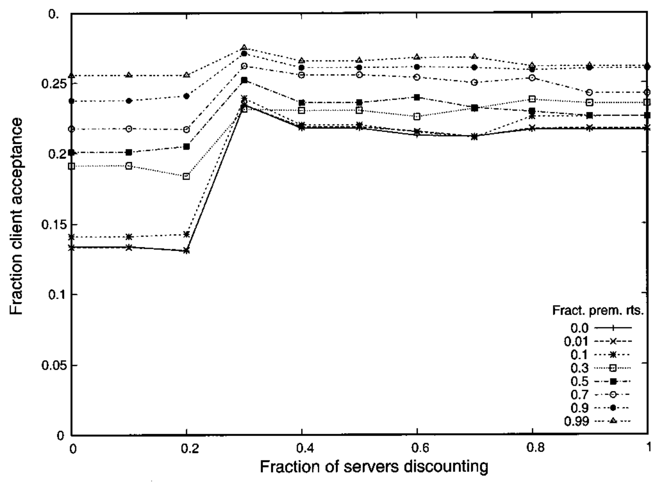

(a)

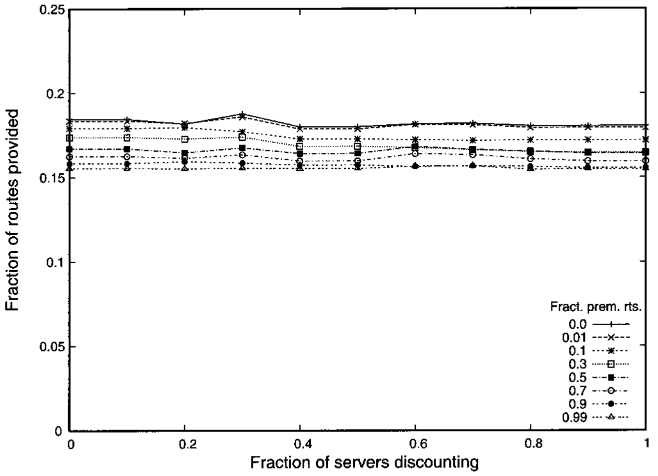

(b)

Fig. 6. AS 1 results with budget (3 6): (a) Acceptance curves and (b) fraction of routes provided.

are two reasons for the apparent lack of coherence between the curves. First, the fraction of client acceptance increase was more dramatic in part because less offers were extended, since AS 1 filled its capacity for connections. Second, more routes were able to be serviced, due to discounts being offered, and this masked the success of AS 1 at capturing route purchases. So, the discounting strategy selected by AS 1 only allowed the AS to maintain its share of routes.

The increase in fraction of routes provided by AS 1, seen on the line with the x markers in Fig. 6(b), had an effect on the overall $P(s)$ for the clients in the network, as shown by the curves in Fig. 4(b). This means that the increase in fraction of routes provided by AS 1 came due to clients getting better service in the form of more routes completed. In fact, the curves show a positive correlation between $P(s)$ and the fraction of servers discounting when there is a low fraction of premium routes.

A larger pattern can be observed from the AS 1 acceptance rate shown in Fig. 6(a). As the fraction of premium clients increases (lines marked 0.01 to 0.99), overall this has a positive effect on the client accept rate for AS 1. This is due to the fact that AS 1 is a tier-1 AS, and its routes are more expensive. As clients process higher numbers of premium route requests (higher bud-

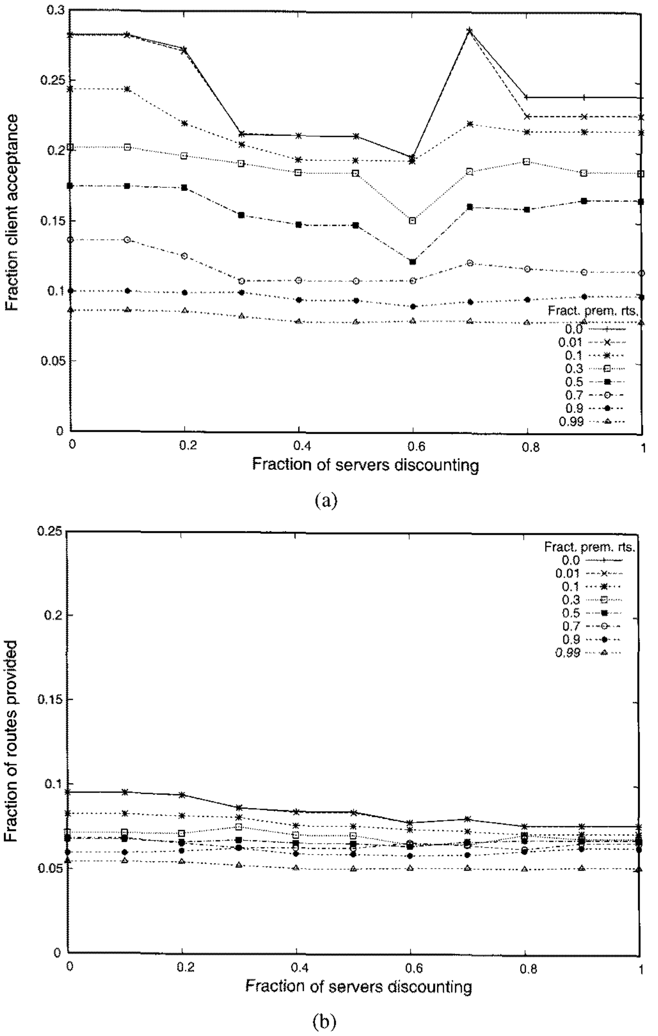

Fig. 7. AS 8342 results with budget (3 6): (a) Acceptance curves and (b) fraction of routes provided.

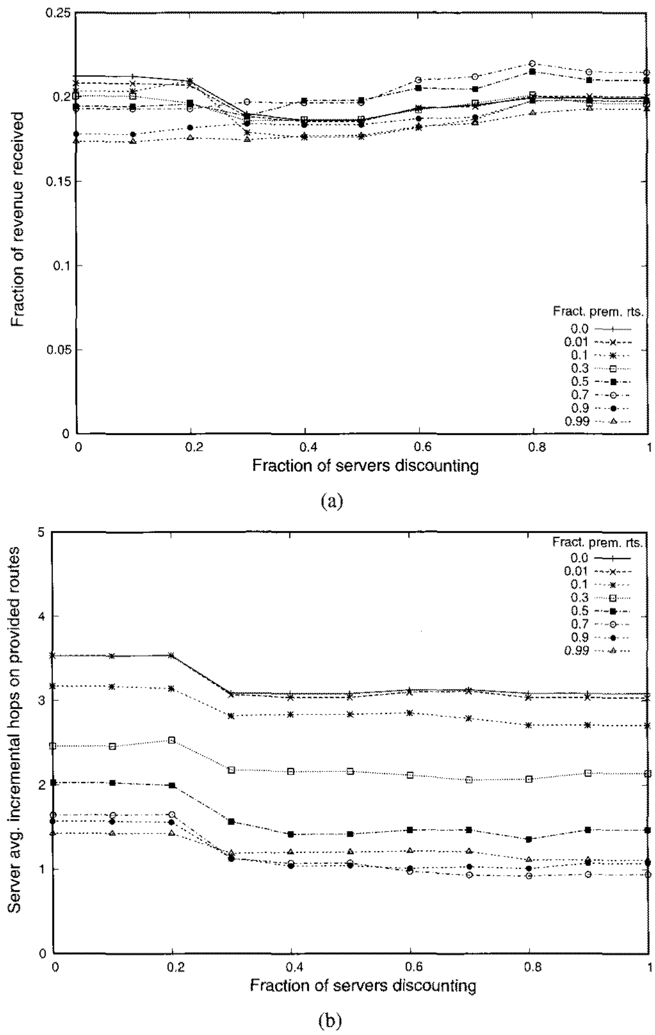

Fig. 8. AS 1 results with budget (3 6): (a) Revenue received curves and (b) average incremental hops.

get and low incremental hop count being more important), AS 1 sees an increase in business. However, the curves trend downwards as higher numbers of ASes discount routes.

Now, we review the same curves for a lower-tier AS, AS 8342. In the fraction of client acceptance curve, Fig. 7(a), the curves exhibit much less stable behavior than for AS 1. For a lower-tier AS, the fraction of client acceptance is more sensitive to the fraction of servers discounting than for a tier-1 AS. The curves have a bathtub shape, due to the movement of clients to other servers who start discounting, but as the cheaper routes become less available, clients return to AS 8342, who offers inherently cheaper routes than the higher-tier servers. In addition, at fraction of servers discounting = 0.7, AS 8342 'decides' to start discounting, which leads to an increase in client acceptance.

The fraction of routes provided curves in Fig. 7(b) show a different shape. When there are less premium routes required, shown by the line with x markers, the lower-tier AS, such as AS 8342, achieves its highest fraction of routes provided, since more clients are handling normal routes that optimize price. Discounting by AS 8342 helps slightly, as some of the curves, labeled 0.3 and 0.5, recover at 0.7 and 0.8 fraction of servers discounting. AS 8342 chooses the price discount strategy at fraction of servers discounting = 0.7.

C. Revenue Share and Incremental Hops

It is important to understand how competition affects revenue for an AS and determine whether or not choosing the strategy of lowering price is effective. This section also reviews the effect of competition on incremental hops of the final route.

Fig. 8(a) shows how AS 1 was able to recapture a much larger fraction of the revenue spent for trusted routes, when it began to discount its routes. The lines marked 0.7 and 0.8 increase dramatically as AS 1 changes its discounting policy from no discounts at fraction of servers discounting = 0.2 to offering discounts at fraction of servers discounting = 0.3. Finally, Fig. 8(b) shows an interesting result occurring over the same interval. The average incremental hops for a trusted path provided by AS 1 decreases from 4.5 down to 4.0 hops on the line marked 0.01. This is the largest change in the incremental hops for AS 1 in the series and it occurs due to a change in discounting policy. Since AS 1 offers routes for less cost, some of its better routes are able to be purchased by clients, who earlier would have accepted less-optimal routes offered by other servers, since they were less expensive.

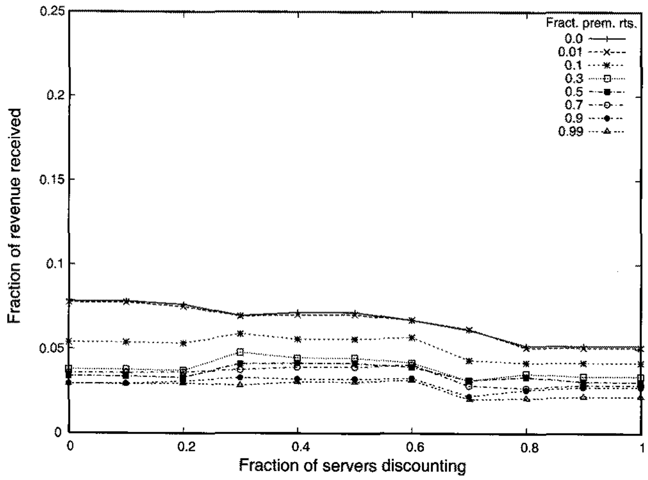

(a)

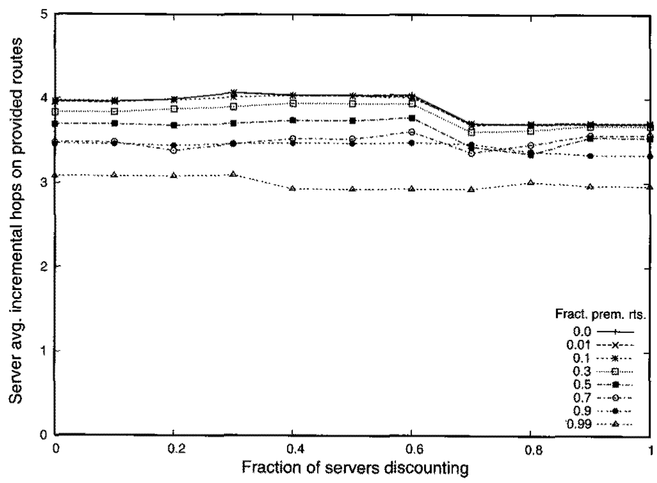

(b)

Fig. 9. AS 8342 results with budget (3 6): (a) Revenue received curves and (b) average incremental hops.

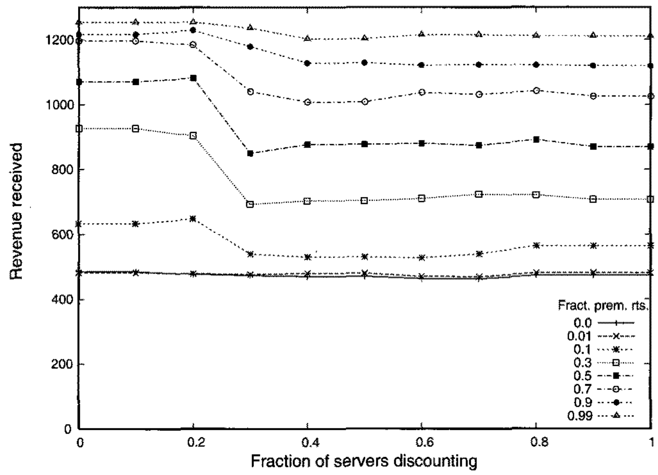

(a)

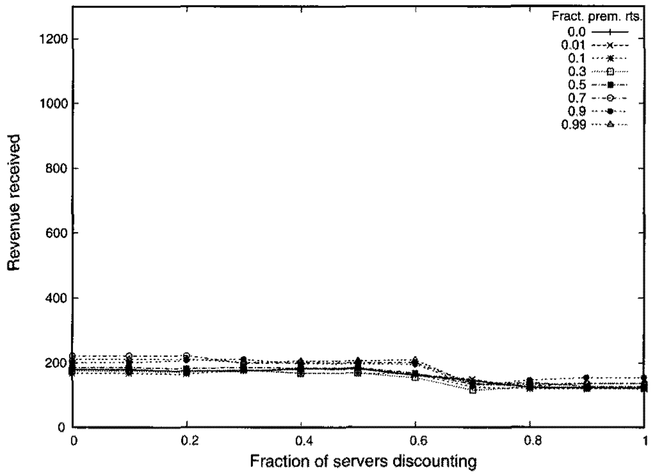

(b)

Fig. 10. Revenue received with budget (3 6): (a) AS 1—tier 1 and (b) AS 8342 - Lower Tier.

Fig. 8(b) shows that, when the fraction of premium routes is high, namely lines labeled 0.9~0.99, the average incremental hops starts out under 1.5 and decreases to 1.1 hops. The client which receives these premium route requests can purchase the low incremental hop count routes offered by AS 1.

The fraction of revenue received for the lower-tier AS, AS 8342, is shown in Fig. 9(a), whose curves follow the same initial trend as those of AS 1. However, there is a very large difference between the curve labeled 0.01 and the other curves. This curve shows that the lower-tier AS is able to bring in a much larger fraction of revenue, when the fraction of premium routes is very low. When the budget is very constrained, clients favor the lower-tier routers which can offer cheaper routes. In addition, the curves show that AS 8342 cannot maintain the same fraction of revenue received as other servers discount.

Fig. 9(b) showing incremental hops for routes provided by AS 8342 reveals that there is a break in all curves when the fraction of servers discounting is between 0.6 and 0.7. This is another effect of AS 8342 deciding to begin discounting in this interval. At this point, clients purchase routing solutions with lower hop count. The incremental hop count curves for AS 8342 demonstrate a range of roughly 1.5 hops, with the best hop count of about 3 hops. In contrast, Fig. 8(b) shows that AS 1 has a larger range of 2 hops, with the lowest curves starting out at about 1.5 hops.

Figs. 10(a) and 10(b) show that neither AS 1 nor AS 8342 is not able to compete and maintain income after servers start discounting routes. Once either server starts discounting their route prices, their income falls off. Fig. 10(a) shows that AS 1 is able to minimize the loss, as the fraction of premium routes increases. In this case, AS 1 is discounting few routes, and competition from other routers has only a slight effect on income. For example, at 0.99 fraction of premium routes and fraction of servers discounting = 0.4, AS 1 lowers the price on 25.6% of the routes purchased, while at 0.01 fraction of premium routes, AS 1 lowers the price on 96.4% of the routes purchased. The lower price allows AS 1 to run at full capacity on connections in both cases. In contrast, at fraction of servers discounting = 0.3, AS 1 is not lowering the price at all, and at 0.01 fraction of premium routes, AS 1 is only running at 85.6% of capacity. Fig. 10(b) shows the flat revenue of the lower-tier AS. In addition, the fraction of premium routes has very little effect on the revenue received by AS 8342.

Figs. 11(a) and 11(b) illustrate the competition that is oc-

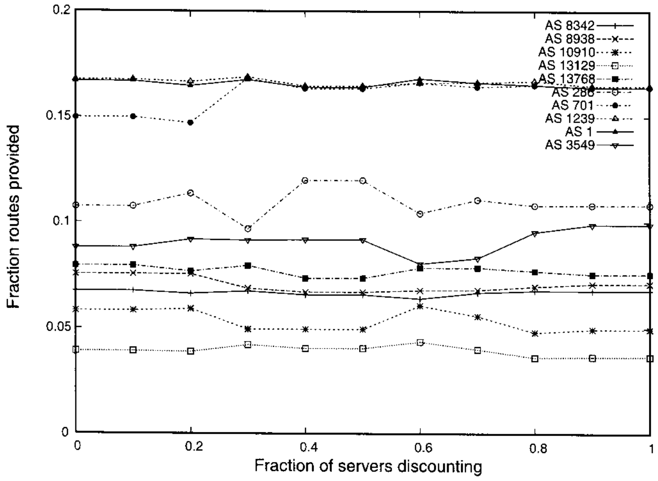

(a)

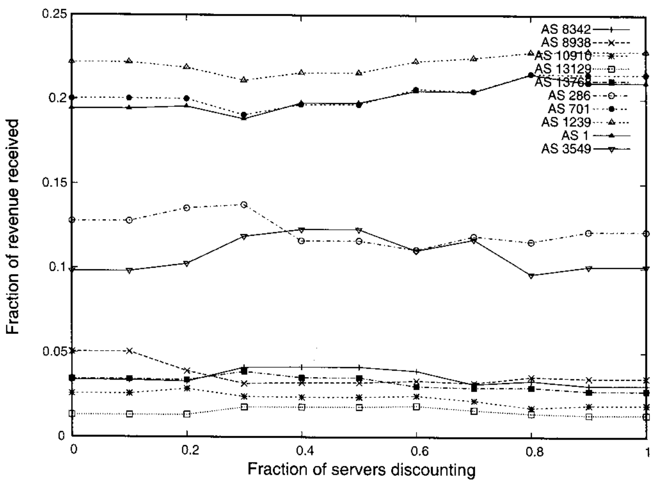

(b)

Fig. 11. Competition among all ASes with budget (3 6): (a) Fraction of routes provided and (b) fraction of revenue received.

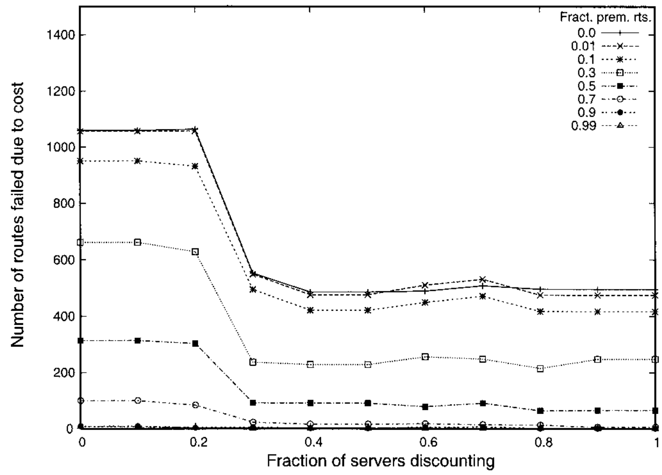

(a)

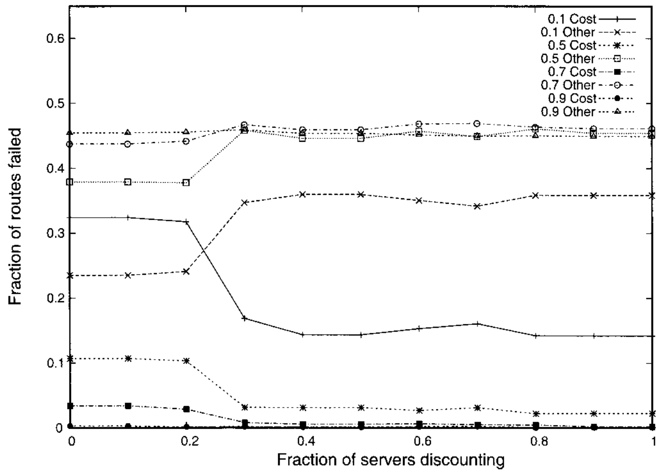

(b)

Fig. 12. Route failure summary with budget (3 6): (a) Routes failing due to cost and (b) fraction of route failure study.

Table 2. Metrics with 0.3 fraction of premium routes and budget (3 6).

Metric	Fraction of servers discounting		
	0.0	0.5	1.0
No response to request	481	868	819
Routes failed due to cost	662	228	248
Routes provided	1456	1503	1532

curring among the ASes for 'market share'. For example, in Fig. 11(a), the lines marked with circles show two ASes which cannibalize each others' clients. As they transition from 0.2 to 0.3 fraction of servers discounting, AS 701 increases the fraction of routes provided, while in the same interval AS 286 decreases the fraction of routes provided by a comparable amount. The routes serviced for one particular client (AS 4651) decreased 61% (from 18 down to 7) for AS 286, while in the same interval, the routes serviced for the client by AS 701 increased 91% (from 21 to 40). This movement of customers is because AS 701 chooses the discounting strategy at fraction of servers discounting = 0.3, and discounts most routes that are purchased. The client AS is directly connected to AS 701, but chooses to utilize AS 286, which is 3 hops away, for routing services, since the routes are more economical. However, when AS 701 discounts the lower-hop-count routing solutions, the customer AS changes servers.

D. Route Request Failures

Examining the cause of client routing request failures as servers compete reveals the pressure that competition has on the network. This pressure on the network causes some clients to be unable to find routing solutions. Figs. 12(a) and 12(b) show causal information about route request failures. Fig. 12(a) shows the number of routes that failed due to cost, that is, when there was a route available, but the client budget would not allow a purchase to be made. As the IDCR servers begin to discount at fraction of servers discounting = 0.3, the number of route failures due to cost falls dramatically. In fact, the curves labeled 0.7 and above decrease almost to 0. This is important, as this is the point where most of the cost-sensitive clients have been served.

Fig. 12(b) shows more detail regarding the failure rate of routes, due to revenue and other causes. To correlate with the prior figure, the line labeled '0.7 cost' falls almost to 0 at a fraction of servers discounting = 0.3. This figure shows that more cost-sensitive clients are being serviced, but what we have discovered is that an increased number of clients receive

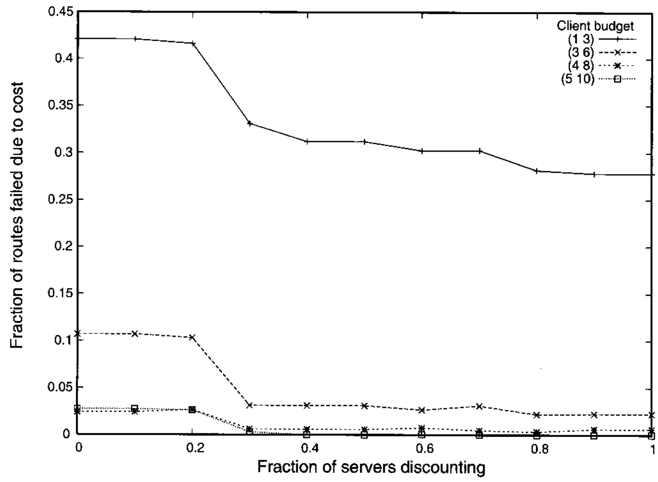

Fig. 13. Routes unrouteable due to route cost at 0.5 fraction premium routes.

no responses to routing requests as more servers discount. Table 2 shows a breakdown of client routing numbers, clearly demonstrating the growth in route failures due to lack of response. When no servers are discounting, there are 481 requests that receive no response, but when all servers discount, this number increases to 819, a 70% increase.

There are two reasons for no response from a server: 1) No route is possible or 2) no resources are available. Since there were only 481 failures when no servers were discounting, this implies that there were at most 481 cases where no routing solution existed, assuming all failures were due to this cause. Therefore, all the new failures that appear are due to lack of resources. As illustrated in Section III, some clients have found that the required server or link is out of resources, leaving them with no routing solution. This implies that the incentive of increasing cost for collaboration with a more heavily-loaded server is not able to move sufficient clients to alternate servers or routes. In addition, further analysis of the routing failures shows that constrained link capacity plays a large role in increasing routing failures for some clients.

E. Budget

Model runs were performed at several other budget value pairs. The graphs of results for these runs are not included in this work. As the budget is changed, the shapes of the curves are similar. Fig. 13 shows how the fraction of routes unrouteable due to cost decreases as the budget is increased. The line with the x markers shows how the failures due to cost fall off quickly above 0.2 fraction of servers discounting. The choice to focus on data for the budget of (3 6) was made due to the fact that, at this cost budget, most routes are economical enough for the clients.

F. Equilibrium with 110-Node Model

We studied the equilibrium of the 110-node model, by running the model with only one independent variable, the fraction of premium routes, but using many different seed values ($n = 100$). The fraction of servers discounting was fixed at 0.5 for all runs. Each server has a 50% probability of choosing the

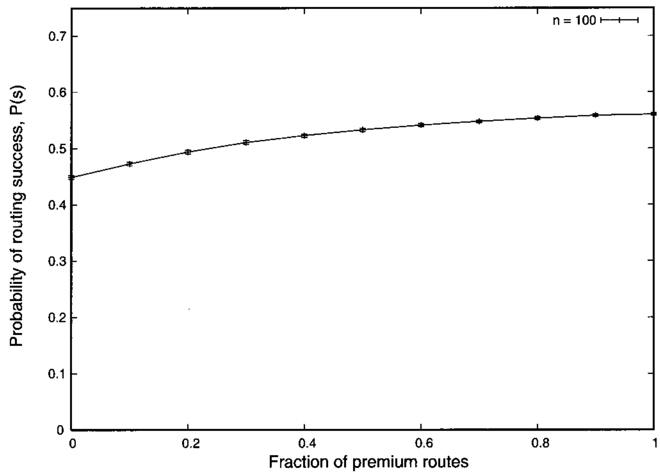

(a)

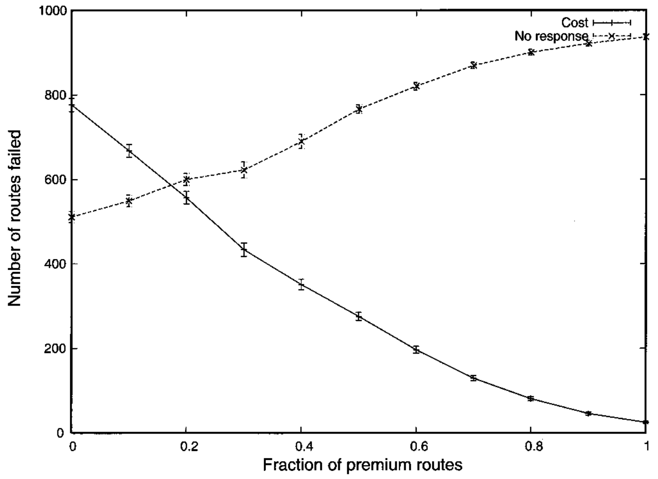

(b)

Fig. 14. Stability study with multiple seeds: (a) Probability of success, $P(s)$ and (b) failed routes.

discounting strategy. The budget was set to (3 6) as in all the prior model runs. Fig. 14(a) shows the gradual increase in the probability of routing success, $P(s)$, as the fraction of premium routes increases. However, Fig. 14(b) shows how the number of clients which do not receive a response to a routing request increases almost 2x. The number of routes failed due to cost decreases even faster, creating the increase in $P(s)$. This is a more complete picture of the issue with no routing response shown in Table 2. An increase in routing success due to more premium routes is expected, but the side effect of increasing routing requests seeing no response is an issue.

G. Server Capacity Sensitivity

The capacity of the servers, although fixed for all prior runs, was varied in a set of runs to illustrate the impact of this dependent variable. In Fig. 15(a), the line labeled 160 shows the effect of too few resources in each server. Even when the fraction of premium routes climbs above 0.2, meaning clients have more money for a route purchase, the clients are unable to purchase the routes. The line is relatively flat. The line labeled 340 shows the other extreme. The servers have plenty of capacity, so as clients are able to afford more routes, $P(s)$ increases in the

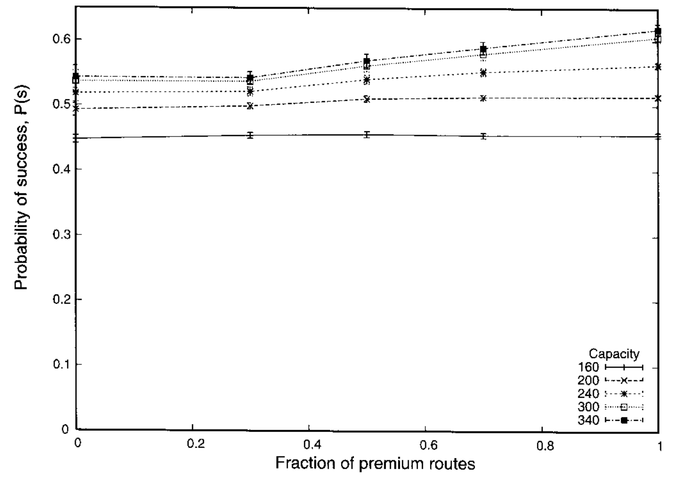

(a)

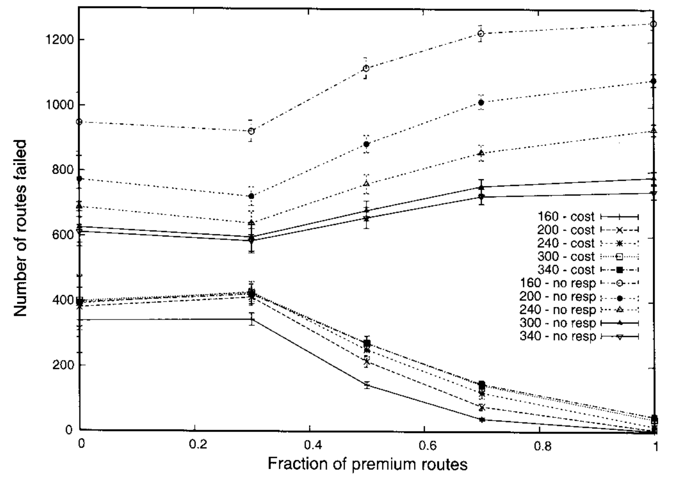

(b)

Fig. 15. Stability study varying server capacity: (a) Probability of success, $P(s)$ and (b) failed routes.

same region of greater than 0.2 fraction of premium routes.

Fig. 15(b) shows the effect of varying server capacity on routing failures. The line labeled "340–cost" shows that above 0.2 fraction of premium routes the number of routes failing due to cost starts lower and drops dramatically. In contrast, the other cost lines start higher, and drop more slowly. Most important are the lines for failures due to lack of response. Note the difference between the lines, labeled "160–no resp" and "340–no resp." When servers have a constrained capacity of 160 connections, there are 55% more failures than with a capacity of 340 at a fraction of premium routes equal to 0. When there are only premium routes being handled, the difference becomes more extreme, with a capacity of 160 resulting in 71% more failures than running with a capacity of 340. This explains why the model run using a capacity of 160 connections sees no increase in $P(s)$. Critical server capacity fills rapidly, and these critical servers are unable to respond to clients which need their collaboration to complete a route.

H. Other Models

Similar curves were generated using the 208-node and 32-node models. The results were similar to those presented for 110-node. For these and additional results, please see [1].

VII. CONCLUSION

We have created several models to evaluate the effects of competition among routers providing an inter-domain routing service, called IDCR. We find that, like a large company in the free market, higher-tier ASes are able to capture a large share of the need for the routing service, and are able to recapture market share, even when lower-tier routers discount the services offered. In addition higher-tier ASes can grow revenue by choosing the discount strategy. Lower-tier ASes cannot maintain market share when higher-tier ASes discount, and cannot maintain revenue.

We have shown that competition among routers, when clients have constrained budgets, has a positive effect on the number of clients who can purchase services with budget constraints. However, the side effect of the competition results in a loss of routing solutions for other clients. These clients see no response from a server when requesting service. Our work shows a large percentage of these failures is due to link overloading.

In the future, we will work on the following areas: handling more than one simultaneous untrusted AS with IDCR, measuring routing overhead of IDCR, and experimenting with strategies to mitigate routing failures due to competition.

REFERENCES

[1] M. Nicholes, "Network security in the core Internet," Ph.D. dissertation, University of California, Davis, 2009.

[2] E. Rasmusen, *Games and Information: an Introduction to Game Theory*. Basil Blackwell Inc., 1989.

[3] F.-N. Pavlidou and K. Koltsidas, "Game theory for routing modeling in communication networks—a survey," *J. Commun. Netw.*, vol. 10, no. 3, pp. 268–287, 2008.

[4] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Trans. Netw.*, vol. 1, no. 5, pp. 510–521, 1993.

[5] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, "Pricing in computer networks: Motivation, formulation, and example," *IEEE/ACM Trans. Netw.*, vol. 1, no. 6, pp. 614–627, 1993.

[6] X. Cao, H. Shen, R. Milito, and P. Wirth, "Internet pricing with a game theoretical approach: Concepts and examples," *IEEE/ACM Trans. Netw.*, vol. 10, no. 2, pp. 208–216, 2002.

[7] R. La and V. Anantharam, "Optimal routing control: Repeated game approach," *IEEE Trans. Autom. Control*, vol. 47, no. 3, pp. 437–450, 2002.

[8] K. Lója, J. Szigeti, and T. Cinkler, "Inter-domain routing in multiprovider optical networks: Game theory and simulations," *Next Generation Internet Netw.*, pp. 157–164, 2005.

[9] L. Li and C. Chen, "Exploring possible strategies for competitions between autonomous systems," in *Proc. IEEE ICC*, 2008.

[10] Y. Xi and E. Yeh, "Pricing, competition, and routing for selfish and strategic nodes in multi-hop relay networks," in *Proc. IEEE INFOCOM*, 2008.

[11] S. Shelford, G. Shoja, and E. Manning, "Achieving optimal revenues in dynamically priced network services with QoS guarantees," *Comput. Netw.*, vol. 51, no. 11, pp. 3294–3304, 2007.

[12] R. D. McKelvey, A. M. McLennan, and T. L. Turocy. (2007). Gambit: Software tools for game theory. Version 0.2007.01.30. [Online]. Available: http://www.gambit-project.org

[13] J. Cowie, A. Ogielski, and D. Nicol. (2002). The SSFNet network simulator. [Online]. Available: http://www.ssfnet.org/homePage.html

[14] B. Premore. (2000). Multi-AS topologies from BGP routing tables. [Online]. Available: http://www.ssfnet.org/Exchange/gallery/asgraph/src.tar.gz

[15] W. Xu and J. Rexford, "MIRO: Multi-path interdomain routing," *ACM SIGCOMM Comput. Commun. Review*, vol. 36, no. 4, pp. 171–182, 2006.

Martin O. Nicholes received his B.S. in Engineering from San Jose State University and his M.S. and Ph.D. in Electrical and Computer Engineering from the University of California, Davis. His major area of focus is computer networking with a minor focus in computer architecture. His research interests include: Network security in the core Internet, and trusted virtual machines. During his 23 years with Hewlett-Packard, he worked in Marketing Technical Support, Manufacturing Test Engineering, and Research and Development. His contributions at Hewlett-Packard were in the areas of server firmware development, server firmware architecture, and server security. He represented Hewlett-Packard in the Trusted Computing Group Industry Standard, as chair of the Server Work Group. He has presented at the Intel Developers Forum and the HP Technical Forum. A collaborative work with HP Labs on virtual machines was selected for presentation at the HP Technical Conference. He currently works for Insyde Software as a Firmware Architect. He holds 10 patents in the area of firmware architecture.

Chen-Nee Chuah is currently a Professor in the Electrical and Computer Engineering Department at the University of California, Davis. She received her B.S. in Electrical Engineering from Rutgers University, and her M.S. and Ph.D. in Electrical Engineering and Computer Sciences from the University of California, Berkeley. Before joining UC Davis in 2002, she spent 9 months as a Visiting Researcher at Sprint Advanced Technology Laboratories. Her research interests lie in the area of computer networks and wireless/mobile computing with emphasis on Internet measurements, network management, anomaly detection, online social networks, vehicular adhoc networks, and multimedia and emerging applications. She received the NSF CAREER Award in 2003, and the Outstanding Junior Faculty Award from the UC Davis College of Engineering in 2004. In 2008, she was named a Chancellor's Fellow of UC Davis. She has served on the Executive/Technical Program Committee of several ACM and IEEE conferences and is currently an Associate Editor for IEEE/ACM Transactions on Networking.

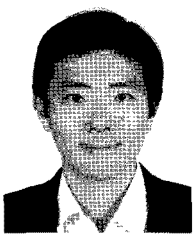

Felix Wu is currently a Professor of Computer Science at UC Davis. He has been doing 'experimental' system research, i.e., building prototype systems to justify and validate novel architectural concepts. Since 1995, he and his students have built many experimental systems in the areas of fault tolerant network, IPSec/VPN security policy, attack source tracing, wireless network security, intrusion detection and response, and online social network systems. He received his B.S. from Tunghai University, Taiwan, in 1985, and Ph.D. from Columbia University in 1995, all in Computer Science. He has 100+ conference and journal publications.

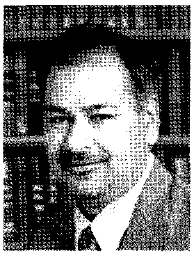

Biswanath Mukherjee holds the Child Family Endowed Chair Professorship at University of California, Davis, where he has been since 1987, and served as Chairman of the Department of Computer Science during 1997 to 2000. He received the B.Tech. (Hons) degree from Indian Institute of Technology, Kharagpur, in 1980, and the Ph.D. degree from University of Washington, Seattle, in 1987. He served as Technical Program Co-Chair of the Optical Fiber Communications (OFC) Conference 2009, and General Co-Chair of OFC 2011. He served as the Technical Program Chair of the IEEE INFOCOM '96 conference. He is Editor of Springer's Optical Networks Book Series. He serves or has served on the Editorial Boards of eight journals, most notably IEEE/ACM Transactions on Networking and IEEE Network. He is Steering Committee Chair of the IEEE Advanced Networks and Telecom Systems (ANTS) Conference, and served as General Co-Chair of ANTS in 2007 and 2008. He is co-winner of the Optical Networking Symposium Best Paper Awards at the IEEE Globecom 2007 and IEEE Globecom 2008 conferences. He won the 2004 UC Davis Distinguished Graduate Mentoring Award. He also won the 2009 UC Davis College of Engineering Outstanding Senior Faculty Award. To date, he has supervised to completion the PhD Dissertations of 46 students, and he is currently supervising approximately 20 Ph.D. students and research scholars. He is author of the textbook "Optical WDM Networks" published by Springer in January 2006. He served a 5-year term as a Founding Member of the Board of Directors of IPLocks, Inc., a Silicon Valley startup company. He has served on the Technical Advisory Board of a number of startup companies in networking, most recently Teknovus (acquired by Broadcom), Intelligent Fiber Optic Systems, and LookAhead Decisions Inc. (LDI). He is a Fellow of the IEEE.