



2010

Controlling the Growth of Internet Routing Tables Through Market Mechanisms

Kin-Wah Kwong

University of Pennsylvania, kkw@seas.upenn.edu

Roch Guerin

University of Pennsylvania, guerin@acm.org

Follow this and additional works at: http://repository.upenn.edu/ease_papers

Recommended Citation

Kin-Wah Kwong and Roch Guerin, "Controlling the Growth of Internet Routing Tables Through Market Mechanisms", . January 2010.

Suggested Citation:

Kwong, Kin-Wah and R. Guerin. "Controlling the Growth of Internet Routing Tables Through Market Mechanisms." In *Proc. 2010 ACM Workshop on Re-Architecting the Internet*. (ReArch'10), Philadelphia, PA, November 2010.

© ACM, 2010. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Proc. 2010 ACM Workshop on Re-Architecting the Internet* (ReArch'10.)

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/ease_papers/561
For more information, please contact libraryrepository@pobox.upenn.edu.

Controlling the Growth of Internet Routing Tables Through Market Mechanisms

Abstract

The growth of core Internet routing tables has been such that it is now viewed as an impediment to the continued expansion of the Internet. The main culprit is multi-homing that stems from sites' desire for greater reliability and diversity in connectivity. These locally rational decisions have a global impact on the Internet, and there is currently no mechanism to effectively control them. A number of technical solutions are being pursued, but this paper explores the use of a "market mechanism." It formulates a model that accounts for sites' incentives and the impact their connectivity choices have on the size of routing tables, and introduces a pricing scheme that seeks to better reapportion the resulting costs. The model is solved for two configurations that capture different deployment realizations and stages. They demonstrate the scheme's effectiveness in controlling the growth of Internet routing tables, while improving the welfare of sites and Internet Service Providers.

Keywords

pricing, multi-homing, market mechanisms, routing table growth

Comments

Suggested Citation:

Kwong, Kin-Wah and R. Guerin. "Controlling the Growth of Internet Routing Tables Through Market Mechanisms." In *Proc. 2010 ACM Workshop on Re-Architecting the Internet*. (ReArch'10), Philadelphia, PA, November 2010.

© ACM, 2010. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Proc. 2010 ACM Workshop on Re-Architecting the Internet* (ReArch'10.)

Controlling the Growth of Internet Routing Tables Through Market Mechanisms*

Kin-Wah Kwong and Roch Guérin
University of Pennsylvania
{kkw@seas, guerin@ee}.upenn.edu

ABSTRACT

The growth of core Internet routing tables has been such that it is now viewed as an impediment to the continued expansion of the Internet. The main culprit is multi-homing that stems from sites' desire for greater reliability and diversity in connectivity. These locally rational decisions have a global impact on the Internet, and there is currently no mechanism to effectively control them. A number of technical solutions are being pursued, but this paper explores the use of a "market mechanism." It formulates a model that accounts for sites' incentives and the impact their connectivity choices have on the size of routing tables, and introduces a pricing scheme that seeks to better reappropriate the resulting costs. The model is solved for two configurations that capture different deployment realizations and stages. They demonstrate the scheme's effectiveness in controlling the growth of Internet routing tables, while improving the welfare of sites and Internet Service Providers.

1. INTRODUCTION

The growth of core Internet routing tables has been proceeding at a pace that far exceeds that of the Internet growth, and this is threatening the scalability of the Internet routing system, both in the short-term (most routers have hard limits on the number of routes they can store) and in the long-term (it outpaces the rate of improvement from technology itself). One of the main culprits behind this trend is the desire for individual "sites" to improve the reliability of their Internet connectivity. This desire for better reliability manifests itself primarily through multi-homing decisions [6], *i.e.*, connectivity to the Internet through more than one provider to enjoy the benefits of greater path diversity, and herein lies the problem.

Specifically, because Internet addressing is coupled to address location, site addresses are commonly associated with one of their providers, and consist of a "sub-block" of the provider's original address block, *e.g.*, a site may be allocated a "/20" subnet out of the provider's own "/8" block¹. In the absence of multi-homing, the provider advertises only its own aggregate block, *i.e.*, its "/8" block, and not those of its individual customer sites, so that a single route entry in core Internet routing tables is sufficient to advertise reachability to all sites connected to the provider². In contrast, when sites are multi-homed, their individual address sub-blocks are now

advertised by the different providers to which they connect. This translates into additional entries in core Internet routing tables.

The challenge in tackling this problem is that it stems from locally rational decisions (improving one's connectivity reliability) that have globally harmful effects (the unimpeded growth of core routing tables will ultimately affect all Internet users). In other words, the lack of coupling between individual site decisions and their global impact is in part responsible for the growth of core Internet routing tables. A natural approach is, therefore, to introduce such a coupling and the goal of the paper is to explore if this can be effectively realized through simple market mechanisms.

This paper develops a model that accounts for the cost of larger routing tables that Internet Service Providers (ISPs) incur, and the benefits to sites of the greater reliability that multi-homing affords. The coupling between the two is provided through a *pricing* scheme that charges sites not only for the traffic they generate and receive, but also for the extent to which they contribute entries to core Internet routing tables, *i.e.*, by introducing a per route charge. This represents a departure from current Internet pricing models that only charge based on traffic volume, and is intended to explicitly account for the cost (to the global Internet) of site decisions (*i.e.*, multi-homing or not) that affect routing table sizes. In the model, ISPs set traffic and route charges so as to maximize their profits. Sites then choose connectivity (*e.g.*, multi-homing or not) to in turn maximize their own utility based on (traffic and route) costs and the level of reliability of the different connectivity options available.

Because a general solution to even the simplified Internet model considered in the paper is essentially intractable, we focus on two special cases that capture different deployment scenarios and stages. The first case seeks to assess the eventual benefits to the Internet from a pricing that incorporates a route charge. This is realized through a scenario where all (tier-1) ISPs have adopted such a scheme, and comparing its outcome to that of a scheme where ISPs continue charging sites based only on traffic volume. The second case explores a scenario aimed at emulating incremental deployment. Specifically, a single tier-1 ISP introduces a pricing scheme that includes both route and traffic charges, while other tier-1 ISPs stay with a pricing that charges only for traffic. The goal of this scenario is to assess both the level of success of the ISP that introduced route charges, and the overall effect this has on the Internet, *e.g.*, site welfare.

The paper is organized as follows. Related works are reviewed in Section 2, and Section 3 introduces our model. Section 4 outlines the two scenarios and the approach for solving them, while their analysis is in Section 5. Section 6 summarizes the paper's findings.

*This work was supported by NSF grant CNS-0915982.

¹A "/ x " block refers to a network address identified through its first x (out of 32) bits, so that the network spans a total of 2^{32-x} individual host addresses.

²This assumes the use of "provider aggregatable (PA)" addresses, as opposed to "provider independent (PI)" addresses. The latter are a minority [7] and for simplicity ignored in this paper.

2. RELATED WORKS

A number of recent works address the growth of core Internet routing tables. One approach is to eliminate some of the factors contributing to this growth by decoupling addressing from location information, *e.g.*, CRIO [11] and LISP [1]. Those works rely on a mapping from end-system addresses (identifiers) to locators, with locators being used for routing purposes. The decoupling of end-system addresses from their locations removes one of the underlying causes that resulted in an increase in routing table entries for addresses accessible from multiple locations. Another approach, *e.g.*, [10], is to rewrite local addresses at the edges using transit addresses that are more easily aggregated; hence limiting the growth of core routing tables. Although those approaches can be effective, some of the scaling issues are now transferred onto other systems than the routing system.

In contrast to the technical solutions we just alluded to, this work does not seek to change the current routing system. Instead, we want to explore how to better apportion through market mechanisms, the costs to the global Internet (from larger core routing tables) of individual site connectivity decisions. The concept of charging for routes was also proposed in [8], with, however, a different focus and approach. In particular, [8] sought to establish bilateral contractual agreements between network entities so as to control route advertisements and traffic exchange.

The market approach proposed in this paper is only one of many recent proposals that reflect the growing need to account for economic factors in the Internet (see [5] for a recent review). Although none of these recent works focused on controlling the growth of routing tables, many sought to improve various aspects of Internet routing. For example, [3, 4] target the use of market mechanisms in Internet peering agreements to better control the volume of traffic exchanged. Similarly, MINT [9] looks at improving Internet routing by enabling end-to-end paths of greater efficiency.

3. MODEL

We develop a model to explore the use of market mechanisms to control the growth of core Internet routing tables. This includes sites and ISPs, accounts for their respective utility functions, and offers a basic representation of how routes propagate in the Internet. Sites derive revenue from connecting to the Internet, and this revenue is affected by the reliability of their Internet connectivity. Sites are heterogeneous in their revenue, and thus in how they value connectivity reliability. This and the cost of connectivity options determine how they choose to connect to the Internet. Connectivity costs are set by ISPs; typically in the form of traffic charges. In the model, we consider the addition of charges that reflect a site’s contribution to the growth of core routing tables, *i.e.*, based on the number of entries it contributes. The goal is to deter sites from choosing connectivity options that fuel the growth of Internet routing tables, when they derive little added value from them.

We abstract the complex nature of the Internet to a simplified three-tier model that captures its hierarchical nature and basic relationships. Our model consists of sites, tier-2, and tier-1 providers. Sites connect to one or more tier-2 providers, which in turn connect to one or more tier-1 providers to which they advertise their local (customer) routes. Single-homed sites do not contribute additional routes beyond that of their tier-2 provider’s block, while multi-homed sites typically result in each of their tier-2 providers advertising (to their tier-1 providers) a dedicated route entry for the site³. Tier-1 providers are fully meshed, *i.e.*, they peer with and

³The tier-2 provider that allocated the site its addresses does not have to. However, not advertising its own copy of the site’s route

advertise all their customer routes to each other. Tier-1 providers represent the backbone of the Internet and provide global connectivity by maintaining route entries for all connected destinations. The metric of interest in our study is the number of routes they carry. Tier-2 providers connect to at least one tier-1 provider to have global connectivity. We ignore possible (peering) connections between tier-2 providers as they do not affect the size of tier-1 routing tables. Furthermore, as we shall see, tier-2 providers are not decision makers in the model (only sites and tier-1 providers are), and are only present to allow different connectivity options. For simplicity we limit these to dual-homing, *i.e.*, a site connects to one or two tier-2 providers. Similarly, a tier-2 provider connects to up to two tier-1 providers.




Figure 1: Connectivity options for sites.

The connectivity options available to a site are shown in Fig. 1, and are denoted by their “shape,” *i.e.*, ‘I’, ‘Y’, ‘D’ and ‘U’. Because reliability is a key factor in a site’s choice of a connectivity option, we discuss next the reliability of the options of Fig. 1. Connectivity through tier-1 and tier-2 providers is assumed to fail with probabilities p_1 and p_2 , respectively. For simplicity and to reflect typical reliability levels, we assume $0 < p_1 \leq p_2 < \frac{1}{2}$, $p_1 p_2 \approx 0$, and $p_2 p_2 \approx 0$ (quadratic terms are negligible). The failure probability p_C of each of the options of Fig. 1 is then:

$$p_C = \begin{cases} p_1 + p_2 & \text{for option I} & (1a) \\ p_2 & \text{for option Y} & (1b) \\ p_1 & \text{for option D} & (1c) \\ 0 & \text{for option U} & (1d) \end{cases}$$

Note that the zero failure probability of option ‘U’ is why an option involving dual-homing to dual-homed tier-2 providers is absent from Fig. 1, *i.e.*, it does not improve reliability.

Next, we discuss how the options of Fig. 1 affect the propagation of routes. Consider options ‘I’ and ‘Y’ that correspond to single-homing. The site’s addresses (route) are allocated from one of its tier-2 providers’ blocks⁴. As a result, neither option contributes an additional entry to the routing tables of tier-1 providers. Note that because option ‘Y’ has a higher reliability than option ‘I’, it is typically preferred by sites unless it has a higher cost. Option ‘D’ corresponds to a dual-homed site whose tier-2 providers are single-homed to the *same* tier-1 provider. The tier-2 providers advertise a separate route for the site, but it need not propagate beyond the tier-1 provider to which they connect, *i.e.*, it is not present in the routing tables of other tier-1 providers. Advertising the aggregate from which the site’s addresses were allocated (by one of the two tier-2 providers) is sufficient (note that this requires some configuration in the tier-1 provider). Finally, option ‘U’ corresponds to a dual-homed site connected to tier-2 providers that are themselves

would result in all traffic destined to the site flowing through the other tier-2 provider, which could be detrimental to revenue.

⁴Recall that we assume the use of PA addresses.

connected to *different* tier-1 providers. In this case, the site’s individual route propagates to all tier-1 providers.

Next, we discuss the pricing scheme of providers and the prices of connectivity options. As mentioned earlier, tier-2 providers are not decision makers in the model. Hence, we assume that they simply pass on to their customer sites charges levied by tier-1 providers. These have two components; traffic and route charges. Traffic charges are per unit of carried traffic, while route charges apply to site routes. For example, tier-1 providers may equate site routes with routes smaller than, say, a “/16” address block, and charge their tier-2 customers for each such route they advertise. Those costs would then be passed-on by the tier-2 providers to the sites associated with the routes. Tier-1 providers may also charge each other for site routes, *i.e.*, smaller than a “/16” address block, although such charges are likely to depend on bilateral agreements and vary across providers. We discuss in Sections 4.1 and 4.2 two illustrative examples for how route charges could arise.

Route and traffic charges influence how sites choose connectivity options. Sites have a utility function and select the option that maximizes it. The utility $h_s(C)$ of site s under option C is defined as

$$h_s(C) = m_s(1 - p_C) - T_C - R_C \quad (2)$$

where m_s denotes the site’s revenue, $m_s(1 - p_C)$ is the site’s effective revenue given the reliability of option C , and T_C and R_C are the traffic and route charges under option C . Under the model of $h_s(C)$, a site opts for a more reliable option only if its improvement in revenue exceeds its increase in (traffic and route) costs. To simplify the analysis, we assume that sites’ revenues are uniformly distributed in $[0, 1]$. Although sites exhibit heterogeneity in their revenue, they are assumed homogeneous in their (bi-directional) traffic, which without loss of generality is taken to be one unit. $h_s(C)$ defines how a site selects its Internet connectivity *once* R_C and T_C are known. These are in turn set by the tier-1 providers that choose them so as to maximize their own utility, given the site’s decision process. Specifically, the utility $U_i(\alpha_i, \beta_i)$ of tier-1 provider i that charges α_i for each copy of a site route it receives, and β_i for each unit of traffic it carries is of the form

$$U_i(\alpha_i, \beta_i) = \left(n_I^{(i)} + n_D^{(i)} + \frac{n_Y^{(i)} + n_U^{(i)}}{2} \right) \beta_i + \left(\hat{n}_D^{(i)} + \hat{n}_U^{(i)} \right) \alpha_i - \left(n_D^{(i)} + n_U \right) b. \quad (3)$$

In $U_i(\alpha_i, \beta_i)$, the first term corresponds to the traffic revenue the tier-1 provider receives, where we have assumed that the (bidirectional) traffic from sites connected to two tier-1 providers (options ‘Y’ and ‘U’) is split equally between them; hence the factor $\frac{1}{2}$. $n_C^{(i)}$ denotes the number of sites picking option C via provider i . The second term deserves additional explanations. It tracks the tier-1 provider’s revenue from route charges. These include revenues generated from charges to the provider’s multi-homed customers and possibly charges to other tier-1 providers that advertise routes from their own multi-homed customers. Conversely, these other tier-1 providers may charge provider i whenever it advertises the site routes of its multi-homed customers. Those charges must be subtracted from the provider’s revenue. Positive and negative charges are combined in the terms $\hat{n}_D^{(i)}$ and $\hat{n}_U^{(i)}$ that track the aggregate contributions to provider i from sites that selected options ‘D’ and ‘U’ respectively. In Section 4, we discuss two representative scenarios of $\hat{n}_D^{(i)}$ and $\hat{n}_U^{(i)}$. Finally, the last term captures the factor that motivated this study in the first place, namely the fact that there is a cost ($b \geq 0$ is a constant) associated with route entries, *e.g.*, from more expensive routers or additional high-speed




Figure 2: Sample site utility across connectivity options.

memory in existing routers. This factor focuses on the total number of site routes ($n_D^{(i)} + n_U$) in the Internet, as they are the dominant contributor to the growth of routing tables, and for analytical tractability is assumed linear in this number. The next section discusses solutions to the provider’s utility maximization problem in two special scenarios.

4. SOLUTION METHODOLOGY

In principle, given the route propagation framework of Fig. 1, the traffic and route costs, T_C and R_C , of each connectivity option can be computed given the values of route and traffic charges, α_i and β_i , that tier-1 providers charge. As a result, given α_i ’s and β_i ’s, tier-1 providers can predict how many sites choose which connectivity option. A representative example of site utility across connectivity options is shown in Fig. 2 as a function of site revenue, together with the resulting site decisions (thick dashed line) and thus number of sites choosing each option. Note that changing α_i ’s and β_i ’s, and hence R_C and T_C , shifts the intersections of the different utility functions. This not only affects the number of sites that choose a particular option, but can also altogether eliminate some from consideration, *i.e.*, as they are strictly dominated by others in the range $m_s \in [0, 1]$.

With the number of sites choosing each connectivity option in hand, providers can compute the values of the different variables in $U_i(\alpha_i, \beta_i)$, and therefore their own utility. Unfortunately, as shown in Fig. 2, although site utility is continuous, its derivative experiences discontinuities at option boundaries. Finding optimal traffic and route charges must, therefore, be done by evaluating $U_i(\alpha_i, \beta_i)$ for all possible combinations of connectivity options that can arise as α_i and β_i , (R_C and T_C) vary. We discuss this procedure in more details below in the context of two special scenarios.

The first scenario is one where all tier-1 providers have agreed to levy route charges, including settlements for such costs between themselves. This captures a setting where route charges have been widely adopted across the Internet, and provides insight into the impact this would have on Internet growth and the growth of routing tables. To focus the investigation on this core aspect, route and traffic charges are assumed homogeneous across providers, and equal to α and β , respectively. The second scenario acknowledges that systematic adoption of route charges is unlikely, at least at first. It instead assumes that a single tier-1 provider proceeds to introduce route charges, and that it is able to levy them only on its own customers. The goal in this scenario is to investigate the effect of the new provider on site connectivity choices, as well as assess its eventual success or demise.




Figure 3: Connectivity options under ubiquitous route charges.

4.1 Ubiquitous Route Charges

Because all tier-1 providers use route and traffic charges of α and β , we limit the model to two tier-1 providers as this is sufficient to realize the connectivity options of Fig. 1. Note that under these assumptions, option ‘Y’ always yields a higher site utility than option ‘I’ because of its greater reliability. As a result, the only options that need to be considered in this scenario are ‘Y’, ‘D’, and ‘U’ as shown in Fig. 3, and their $h_s(C)$ expressions are:

$$h_s(C) = \begin{cases} m_s(1 - p_2) - \beta & C = Y \quad (4a) \\ m_s(1 - p_1) - \beta - 2\alpha & C = D_{1/2} \quad (4b) \\ m_s - \beta - 4\alpha & C = U \quad (4c) \end{cases}$$

The term -2α in the site utility of option ‘D’ comes from the route charges that the tier-1 provider levies from the two tier-2 providers that advertise the site route. As mentioned earlier, those charges are passed on to the site by the tier-2 providers. Similarly, the term -4α in the site utility of option ‘U’ comes from two components, namely the route charges (2α) again levied from the two tier-2 providers, and two additional route charges introduced by the settlement of route charges between the two tier-1 providers. These charges are again passed on to the sites, and when compared to option ‘D’, capture the broader propagation of the route (to all tier-1 providers).

The utility of either tier-1 provider is easily found to be of the form

$$U(\alpha, \beta) = \left(\frac{n_Y + n_D + n_U}{2} \right) \beta + \left(\frac{n_D}{2} + n_U \right) (2\alpha - b). \quad (5)$$

Note that the factor $\frac{1}{2}$ applied to the n_D sites’ contributions to traffic and route revenues is because sites that select this option are assumed equally split across both tier-1 providers.

4.2 Incremental Deployment Scenario

In this scenario, a single tier-1 provider has opted to charge its (tier-2) customers for site routes they advertise. Because those charges are in addition to traffic charges, it is clear that to attract customers, the tier-1 provider has to offer lower traffic charges than other tier-1 providers. More specifically, we assume that other tier-1 providers charge β per unit of traffic, but do not charge for routes. For simplicity and because it is sufficient to accommodate the connectivity options of Fig. 1, we limit the number of existing providers to 2 (labeled as 1 and 2). The tier-1 provider that introduces route charges is numbered provider 3. It charges α for site routes and $\beta_3 < \beta$ per unit of traffic. We further assume that tier-1 providers 1 and 2 have selected β so as to optimize their revenue given existing costs⁵, and do not revisit their decision after provider 3 enters the market. This is representative of the situation shortly after the emergence of provider 3.

⁵In the absence of route charges and barring other penalty for multi-homing, it can be easily shown that optimal $\beta = \frac{1}{2} + b$.




Figure 4: Connectivity options under incremental deployment scenario.

In order to understand how provider 3 would set its route and traffic charges, it is necessary to identify the different connectivity options that sites will then have and what site utility is associated with each of them. Site connectivity options in this scenario are shown in Fig. 4, which displays the greater complexity of this scenario because of the larger number of possible combinations it gives rise to. Specifically, sites with low revenue will choose option ‘I’ since it has the lowest traffic costs (recall that $\beta_3 < \beta$) and no route costs. Option ‘Y’ has slightly higher traffic costs (half of the traffic is routed through either provider 1 or 2), but has higher reliability and still no route costs. Option ‘D’ incurs route costs, but sees further reliability improvements and low traffic costs, so that it may be a viable option. Option ‘U’_{1,2} has high traffic costs but no route costs and the highest reliability, while option ‘U’_{1/2,3} boasts the same reliability and lower traffic costs (half of the traffic is through provider 3) but incurs route costs. The $h_s(C)$ expressions for different options are

$$h_s(C) = \begin{cases} m_s(1 - 2p) - \beta_3 & C = I \quad (6a) \\ m_s(1 - p) - \frac{\beta + \beta_3}{2} & C = Y \quad (6b) \\ m_s(1 - p) - \beta_3 - 2\alpha & C = D \quad (6c) \\ m_s - \beta & C = U_{1,2} \quad (6d) \\ m_s - \frac{\beta + \beta_3}{2} - \alpha & C = U_{1/2,3} \quad (6e) \end{cases}$$

where in order to keep the analysis tractable we assume that $p_1 = p_2 = p$, i.e., tier-1 and tier-2 providers have the same reliability (note that this implies that options ‘Y’ and ‘D’ then have the same reliability). Note the presence of a single route charge (α) for option ‘U’_{1/2,3} versus a charge of 2α for option ‘D’, because provider 3 is the only one levying route charges, so that unlike the scenario of Section 4.1 there is no additional route charges from the other tier-1 provider to propagate down to sites.

Using $h_s(C)$, the utility of provider 3 can be expressed as a function of the number of sites that choose each connectivity option for a given choice of α and β_3 . This corresponds to

$$U_3(\alpha, \beta_3) = \left(n_I + \frac{n_Y}{2} + n_D + \frac{n_{U_{1/2,3}}}{2} \right) \beta_3 + n_D(2\alpha - b) - n_{U_{1,2}}b + n_{U_{1/2,3}}(\alpha - b). \quad (7)$$

As in the case of ubiquitous route charges of Section 4.1, maximizing $U_3(\alpha, \beta_3)$ can be carried out by computing the number of sites that select each of the connectivity options of Fig. 4. This procedure is again outlined in Section 5.

5. ANALYSIS AND DISCUSSION

Solving for the route and traffic charges that maximize tier-1 provider utility in the scenarios of Sections 4.1 and 4.2, calls for

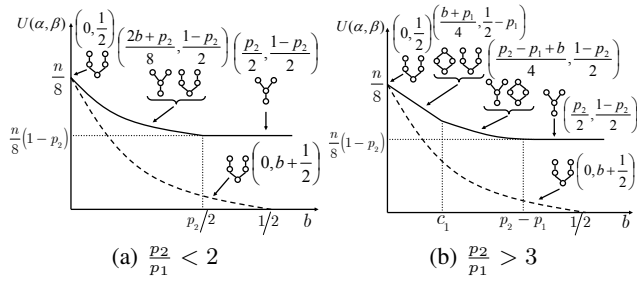


Figure 5: Optimal provider utility and site connectivity options under ubiquitous route charges (solid line) and without route charges (dashed line).

identifying what combinations of connectivity options are meaningful for sites to select as those charges vary.

For example, when route charges are ubiquitous, only options ‘Y’, ‘D’, and ‘U’ need to be considered, and they give rise to at most $2^3 - 1 = 7$ combinations, *i.e.*, a single connectivity option, all possible pairs of options, or all three options, although not all combinations may arise as discussed in Section 4.1. This depends in part on how the x -intercept of the utility functions of the different options varies with α and β , as shown in Fig. 2. Optimizing the provider utility then calls for identifying the optimal (α, β) values for each feasible combination of connectivity options, and ultimately selecting the one that yields the highest utility. Next, we outline the results of this methodology and their implications in the context of the two scenarios of Sections 4.1 and 4.2.

5.1 Ubiquitous Route Charges

As mentioned above, in this scenario there are up to 7 possible combinations of feasible connectivity options that can arise as α and β vary. Each combination is feasible over a certain range of α and β values, and these ranges depend in turn on the relative values of p_1 and p_2 . Details can be found in [2], and this section provides two representative examples of the outcome of the model, *i.e.*, a case where tier-2 reliability is close to that of tier-1 ($p_1 \leq p_2 < 2p_1$ in Fig. 5(a)), and one where it is worse ($p_2 > 3p_1$ in Fig. 5(b)). In both cases, sites select from different subsets of possible connectivity options, with which subset to select from being itself a function of b , the cost of route entries to providers. Note that the value of c_1 in Fig. 5(b) is constant in terms of p_1 and p_2 .

Figs. 5(a) and 5(b) display, as a function of b , the optimal provider utility with and without route charges, the corresponding optimal route and traffic charges α and β , and the connectivity options that coexist under those conditions. In both figures, the two utility curves clearly illustrate the benefits, at least to providers, of including route charges.

In both cases, when $b = 0$ (large routing tables are not a cost factor) pricing is such that sites are not penalized for using option ‘U’, which is therefore the one they all select because of its higher reliability. Conversely, when b is large, providers set the route charges high so that no sites can afford them, and they all select option ‘Y’. At intermediate values of b , other combinations become attractive to some sites as a function of their revenue and the corresponding reliability gains. In other words, as the cost b of individual route entries increases, the ability to deter (through route charges) sites that benefit little from selecting connectivity options that contribute to this cost, enables providers to maintain most of their original utility (a decrease by a factor $(1 - p_2)$). In contrast, without route charges providers’ utility steadily declines to 0 as b reaches $1/2$.




Figure 6: Site connectivity options and optimal utility of providers levying route charges (solid line) and not (dashed line).

In practice this means that as routing tables grow (because all sites decide to multi-home) and the cost of route entries correspondingly increases, most sites end-up being priced out of connecting to the Internet (the traffic charges providers need to levy to offset the growing cost of routing tables exceed the sites’ revenue). In other words, even if providers have more customers, they derive lower profits because of the corresponding higher route costs. We expand on this discussion and its implications in Section 5.3.

Interestingly, the benefits of route charges also extend to sites [2]. The providers’ ability to keep traffic charges roughly constant (at about $\frac{n}{2}$), irrespective of route charges, ensures that the number of sites able to afford (some) connectivity also remains roughly constant (at about $\frac{n}{2}$, where n denotes the number of sites seeking Internet connectivity). In addition, site welfare, defined as the sum of the utility of all the sites that can afford Internet connectivity, *i.e.*, for which $h_s(C) > 0$ for some C , also improves when route charges are introduced. This is because the fraction of sites that can afford to connect remains approximately constant, as opposed to declining with the growing cost of routing tables.

5.2 Incremental Deployment Scenario

As shown in Fig. 4, this scenario is more complex because of the greater number of combinations of connectivity options it can involve. Tackling this issue calls again for partitioning possible (α, β_3) pairs into ranges for which sites consider a given set of connectivity options, and solving for optimal values in each range before identifying a global optimum. Fig. 6 displays two representative results of this procedure; one for a case where providers are reasonably reliable ($p \leq 0.1$) and the other for a case where they are not ($p \geq 0.16$). Note that the values of c_4 and c_7 in the figure are constant in terms of p . The figure shows both similarities and minor differences with the ubiquitous route charges scenario.

Fig. 6 confirms the benefits to providers (provider 3 in this instance) of introducing route charges. In both cases, provider 3 succeeds in attracting sites, including to connectivity options that contribute fewer routes to core Internet routing tables, *i.e.*, options ‘I’ and ‘D’. These sites are both customers of the existing providers, and sites previously unable to connect because of the high traffic charges levied by these providers⁶. As a matter of fact, provider 3 is somewhat unrealistically successful in the case of Fig. 6(a), where it eliminates the existing providers as soon as it emerges. This is because the relatively high reliability of providers makes option ‘U’ unattractive as soon as alternatives with lower traffic charges and no/minimum route charges (options ‘I’ and ‘D’) become available. Because existing providers do not modify their traffic charges, sites

⁶Recall that provider 3 offers lower traffic charges ($\beta_3 < \beta$) to compete with existing providers and offset its route charges.




Figure 7: (a) Number of routes in core routing tables. (b) Improvement in provider utility and number of connected sites.

immediately migrate to provider 3. Arguably, more complex interactions can arise in practice, but this shows the core advantage of introducing route charges, *i.e.*, increasing traffic charges is not the only option to offset growing route costs.

As with the scenario of Section 5.1, those benefits also extend to sites albeit with some minor differences. In particular, not only does overall site welfare improve, the number of connected sites also increases beyond $\frac{n}{2}$ (even when $b = 0$). This is because provider 3 must offer lower traffic charges than existing providers. This constraint is absent in the scenario of Section 5.1, and weakens as b grows which forces existing providers to increase their traffic costs. The latter explains the growth of the utility of provider 3 as b gets larger. Note that the initial decline in Fig. 6(a) is again somewhat artificial. It is because provider 3 needs to set its route and traffic charges to avoid losing customers to the existing providers. This prevents it from increasing route charges to fully offset the increase in route costs from sites selecting option ‘D’.

5.3 Discussion and Extensions

For analytical tractability the paper assumes route costs linear in the number of routes in core routing tables (parameterized by b). In the current Internet $b \approx 0$, and this would remain true if core routing tables were growing at the same pace as technology improves. The concern is that they may not. Capturing this effect is ideally realized through a non-linear function, *e.g.*, quadratic, between route costs and routing table sizes. Such a model unfortunately appears intractable, which motivated the simpler linear model of the paper. In spite of this simplification, the results can be leveraged to explore more complex relationships between route costs and routing table sizes, and we outline next one possible approach.

We mimic a non-linear relation between routing table sizes and their cost by varying b of our linear model over successive periods of Internet growth. Specifically, we assume that at times $t = 0, 1, 2, \dots$, a new set of n sites seeks to connect to the Internet. The route cost b_t at time t is of the form $b_t = a \left(\frac{s_{t-1}}{s_{\max}} \right)^2$, where s_{t-1} is the number of site routes in core routing tables at time $t - 1$, s_{\max} denotes an upper bound on the number of entries that routing tables can accommodate, and a is a normalization factor. The cost of routing tables is initially negligible, *i.e.*, $b_0 = 0$, but as more sites contribute new (site) routes, technology is stressed and route costs increase. In this model, all (tier-1) providers adjust their route and/or traffic charges at each epoch t based on the (route) costs produced by decisions at epoch $t - 1$. Route and traffic charges at epoch t only affect sites seeking to connect at that time, *e.g.*, existing sites are unaffected. The goal is to study the extent to which route charges help control the growth of routing tables, and how they affect the number of sites that connect to the Internet over time.

Consider a representative outcome of Fig. 5(b). Fig. 7(a) shows how route charges curtail the growth of routing tables. Fig. 7(b) shows that this is accomplished while allowing more sites to connect to the Internet, and improving the utility of tier-1 providers.

6. CONCLUSION

This paper reports preliminary results of an investigation of market mechanisms to control the growth of core Internet routing tables. It contemplates adding (site) route charges to the pricing model of ISPs, to better apportion the cost (to the global Internet) of connectivity choices of individual sites. The paper develops a simple model that captures key features of the approach and provides initial evidences of its effectiveness. Extensions to more realistic configurations are the topic of ongoing work.

ACKNOWLEDGMENTS

The authors would like to acknowledge many productive discussions with K. Hosanagar, and the assistance of V. Cherepanov in articulating connectivity choices.

7. REFERENCES

- [1] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis. Locator/ID separation protocol (LISP). IETF Internet-Draft (draft-ietf-lisp-07), April 2010. (Work in Progress).
- [2] K.-W. Kwong and R. Guérin. Controlling the growth of Internet routing tables through market mechanisms. Technical report, University of Pennsylvania, August 2010.
- [3] S. Lippert and G. Spagnolo. Internet peering as a network of relations. *Telecommunications Policy*, 32(1):33–49, 2008.
- [4] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein. Interconnecting eyeballs to content: A Shapley value perspective on ISP peering and settlement. In *Proc. NetEcon*, 2008.
- [5] M. Mueller. Critical resource: An institutional economics of the Internet addressing-routing space. *Telecommunications Policy*, 34(8):405–416, 2010.
- [6] T. Narten. On the scalability of Internet routing. IETF Internet-Draft (draft-narten-radir-problem-statement-05), February 2010. (Work in Progress).
- [7] PI statistics update. RIPE Presentation on Provider Independent Address Space Assignments, October 2006. www.ripe.net/ripe/meetings/ripe-53/presentations/address_space.pdf.
- [8] Y. Rekhter, P. Resnick, and S. Bellovin. Financial incentives for route aggregation and efficient utilization in the Internet. In B. Kahin and J.H. Keller, editors, *Coordinating the Internet*, pages 273–287. MIT Press, 1996.
- [9] V. Valancius, N. Feamster, R. Johari, and V. Vazirani. MINT: A market for Internet transit. In *Proc. ReArch*, 2008.
- [10] C. Vogt. Six/One router: A scalable and backwards compatible solution for provider-independent addressing. In *Proc. MobiArch*, 2008.
- [11] X. Zhang, P. Francis, J. Wang, and K. Yoshida. Scaling IP routing with the core router-integrated overlay. In *Proc. ICNP*, 2006.