

# MINT: A Market for INternet Transit

Vytautas Valancius\*, Nick Feamster\*, Ramesh Johari†, Vijay Vazirani\*  
\*Georgia Tech †Stanford University

## ABSTRACT

Today’s Internet’s routing paths are inefficient with respect to both connectivity and the market for interconnection. The former manifests itself via needlessly long paths, de-peering, etc. The latter arises because of a primitive market structure that results in unfulfilled demand and unused capacity. Today’s networks make *pairwise*, myopic interconnection decisions based on business considerations that may not mirror considerations of the edge networks (or end systems) that would benefit from the existence of a particular interconnection. These bilateral contracts are also complex and difficult to enforce.

This paper proposes MINT, a market structure and routing protocol suite that facilitates the sale and purchase of end-to-end Internet paths. We present MINT’s structure, explain how it improves connectivity and market efficiency, explore the types of connectivity that might be exchanged (vs. today’s “best effort” connectivity), and argue that MINT’s deployment is beneficial to both stub networks and transit providers. We discuss research challenges, including the design both of the protocol that maintains information about connectivity and of the market clearing algorithms. Our preliminary evaluation shows that such a market quickly reaches equilibrium and exhibits price stability.

## 1. Introduction

The Internet’s organic market for connectivity is at a breaking point. The frequency of “de-peering” [5,7,23] that have partitioned the Internet bears witness to the misalignment of incentives between larger ISPs and the edge networks; this tension will become even more apparent as users and applications demand “better than best effort” connectivity. Bandwidth hungry applications are straining the current Internet infrastructure. According to a recent Internet traffic projections report, total IP traffic will increase by a factor of six from 2007 to 2012; strikingly, Video on Demand (VoD) and Voice over IP (VoIP) traffic may increase at annual rates of 44% and 30%, respectively [10]. These applications often have strict bandwidth or delay requirements, and many of them involve interdomain Internet paths (i.e., paths that traverse multiple ISPs).

Unfortunately, two factors—today’s de facto interdomain

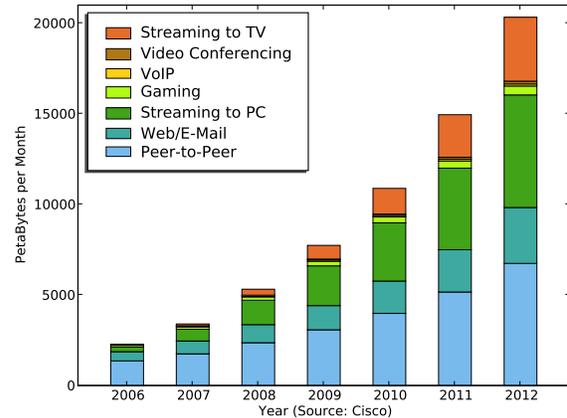


Figure 1: Internet bandwidth growth trends [10].

routing protocol, Border Gateway Protocol (BGP) [27], and the market structure for Internet connectivity—make it difficult for Internet service providers (ISPs) to offer end-to-end connectivity that satisfies specific performance constraints. Various de-peering events (and the resulting disconnectivity) suggest that today’s market structure sometimes lacks incentives for ensuring that end-to-end connectivity exists at all. From the protocol perspective, BGP does not provide any mechanism for ISPs to publicize or exchange performance information about unused capacity, creating inefficiencies where the network may have both excess supply and unfulfilled demand. Indeed, the current network structure and protocols create both connectivity and market inefficiencies by failing to provide information that could be used to match demand with supply.

To account for these shortcomings today, network operators typically resort to reactive, trial-and-error methods to find efficient paths. Most ISPs and edge networks connect to multiple transit networks, and operators continually tune routing configurations in the face of changing network conditions. Such ISPs and stub networks are cornered into purchasing “upstream” connectivity but have little control over the nature or quality of the overall end-to-end path. Government agencies have also recently recognized a need for a more sophisticated market model for Internet capacity [3].

This paper proposes *MINT*, a suite of routing protocols that facilitate a connectivity market for the exchange of *end-to-end* paths between independent networks. Physical connectivity in MINT is the same as it is in today’s Internet. However, instead of (or, perhaps, in addition to) buying upstream connectivity, networks can purchase connectivity to sets of destination networks that satisfy explicit performance

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

ReArch’08 December 9, 2008, Madrid, SPAIN

Copyright 2008 ACM 978-1-60558-234-4/08/0012 ...\$5.00.

metrics, thus providing each network more direct control over some or all of its end-to-end paths. The purchased path can provide direct peering to remote networks; such path is perceived as a direct, one-hop link and, as in the case of the classical peering, it can carry whatever traffic the peers agree to exchange. Such traffic could be service specific (*i.e.* to a specific content servers or VPN customers) as well as general traffic (*i.e.* Internet traffic to the whole region). ISPs providing transit for such paths express their routing policy through prices of path *segments*; they advertise topology using *transit state* updates, each of which contains the price and capacity of connectivity between an ingress-egress pair for that ISP.

Compared to existing routing systems, MINT’s design’s two main salient features are:

- **Policy expression through price.** Unlike BGP, where policies are constrained by bilateral connection agreements between ISPs, path segments in MINT are explicitly priced. MINT’s pricing mechanism subsumes transit policies that are in place today.
- **Ingress-egress pairs as a connectivity abstraction.** Vertices in the transit topology maintained at the transit state database represent ingress and egress points for adjacent ISPs (e.g., an exchange point), and edges represent transit between two such points. Transit updates serve as advertisements for path segments, which form the basis for the connectivity market and inter domain path selection.

MINT resembles “bandwidth brokers” [30] but extends the notion to inter-domain paths. The existing bandwidth markets [18] offer only single-hop reservations; MINT, on the other hand, helps edge networks construct *interdomain* end-to-end paths by composing segments, where each segment may be offered and maintained by a different service provider. Unlike previously proposed micro-payment systems [14], MINT operates at the granularity of long-term path segments, rather than per packet. We believe that a market for transit connectivity more directly reflects the granularity at which networks operate.

MINT’s market for end-to-end paths offers several benefits. First, routing in MINT achieves *greater path diversity* than under conventional routing: there are no “non policy-compliant” paths; every path simply has a price. MINT eliminates the notion of customer-provider and peer-peer links from the Internet topology. By providing explicitly priced end-to-end paths to selected destinations, MINT can avert de-peering. Second, MINT is *flexible*: MINT allows each network to independently select paths to each destination for each type of traffic; routing decisions are not restricted to routes provided by upstream providers as they are with BGP. Third, MINT is *feasible*: it can be implemented using a combination present-day protocols and only minor modifications to the control plane. MINT is also compatible with existing protocols and can operate in parallel with today’s “best effort BGP”.

The design of MINT presents two main challenges: (1) *economic viability* in terms of system stability and incentive structure; and (2) *technical viability* in terms of protocol

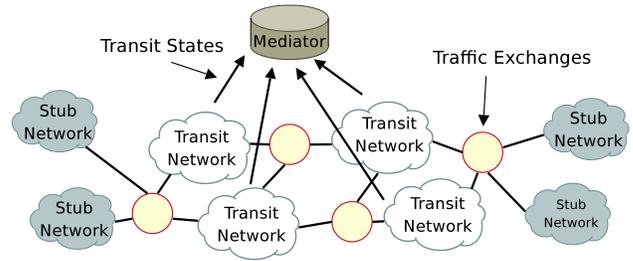


Figure 2: High-level MINT architecture.

scalability as system operates at the Internet scale. To address economic challenges, we devise and simulate an auction and a market structure that allows complementary trading of transit services. Our preliminary evaluation studies convergence of prices and resource utilization; in our ongoing work, we are studying issues relating to technical viability, such as scalability.

In addition to assessing MINT’s economic and technical viability, we must address scaling and privacy concerns. First, MINT potentially propagates a large amount of global state. Second, MINT must provide an *open* connectivity market for all participating domains while preserving the *privacy* of their intra-domain topologies. This paper proposes a high-level system design and operation that is geared to address some of these issues; we also describe how MINT can operate in parallel with existing protocols and without significant changes to the deployed infrastructure.

The rest of the paper is organized as follows. Section 2 provides an overview of MINT. Section 3 outlines MINT’s market operation. Section 4 describes the protocols that implement the market. Section 5 presents a preliminary evaluation showing convergence properties. Section 6 discusses deployment incentives, Section 7 surveys related work, and Section 8 concludes.

## 2. Overview

MINT has three participants: buyers, sellers, and a mediator. The *buyers* may be both large Internet service providers (ISPs) and edge networks. The *sellers* are the large ISPs and, in some cases, edge networks that have more than one inter-domain connection and spare capacity. The marketplace is implemented by a centralized *mediator*, which aggregates the information about the product to be bought and sold—available path segments—and matches demand for end-to-end paths to supply of path segments. Figure 2 illustrates this high-level operation; the market is described in further depth in the following section.

Sellers advertise path segments to the mediator with multiple attributes, including available bandwidth and price. Such advertisements are *transit state updates*; they effectively aggregate information internal to the network and thus preserve privacy of the ISP’s topology. Sellers are encouraged to *under-promise and over-deliver*, where announced information does not necessarily reflect the exact state of the network but rather the state that the ISP is willing to disclose.

To construct an end-to-end path between two autonomous systems, a buyer issues a request with the path’s endpoints

(*i.e.*, the source and destination exchanges), as well as any constraints the path must meet, such as minimum bandwidth, maximum delay, and a list of networks to avoid. Each end-to-end path purchased by a buyer would typically comprise multiple segments; sellers and buyers use the mediator to enable the exchange of goods by matching demand (*i.e.*, end-to-end paths) to supply (*i.e.*, collections of path segments). Note that each network can act as a buyer for some requests (if it acts as a request source) and as a seller for other requests (if it provides transit).

For each request, the mediator finds the lowest cost *exchange-level* path that satisfies the requested constraints. If such a path is found, the mediator returns this path (and the associated resources) to the buyer. For each respective path segment, each network uses distributed path computation to compute and establish *node-level* paths between the network’s ingress and egress.

### 3. Market

This section describes the market operation in MINT. The price of resources in MINT is driven by a *Continuous Double Auction (CDA)*. The CDA is a simple mechanism where buyers and sellers are continuously matched over time. For a single commodity, the market operates as follows. Buyers submit *bids*, consisting of the price they are willing to pay and the quantity they wish to purchase; sellers submit *offers*, consisting of the price they are willing to accept and the quantity they wish to sell. A central mediator matches bids to offers where possible. A bid is only satisfied if offers at the same or lower price are available, with sufficient quantity to match the bid. Both bids and offers are “standing” in the usual CDA: they remain in an *order book* until satisfied or withdrawn. This simple design is employed at stock markets and commodities markets throughout the world. A recent manuscript provides a survey on double auctions [25].

The use of a CDA is complicated by the fact that we are not allocating a single commodity; rather, our market consists of network resources (bandwidth from ISPs between exchange points) that must be *combined* to deliver end-to-end path connectivity. We briefly describe the operation of the CDA we consider in the network setting. As in a standard CDA, offers are “standing”: they do not expire unless explicitly withdrawn by an ISP. Bids, however, *expire immediately*, which simplifies implementation without significantly affecting market operation. Although we focus on an explicitly priced bandwidth market, MINT can be applied as a market implementation tool for any type of interdomain communication services, including best effort connectivity.

As shown in Figure 2, agreements are formed between the mediator and participating ISPs that agree to lease capacity to each other upon request, under the price announced *a priori* to the mediator. These offers can be updated periodically or upon a major network change. Each network participant sets the price of their segments independently. The equilibrium price depends on basic demand and supply laws, as buyers independently make decisions to lease capacity.

Formally, suppose the set of ISPs  $I^{(1)}, I^{(2)} \dots I^{(n)}$  supply bandwidth. The bandwidth available for lease between exchange points  $i$  and  $j$  by ISP  $k$  is advertised to the arbitrator

as a set of offers  $\Pi^{(k)} = \{ \langle P(I_{ij}^{(k)}), B(I_{ij}^{(k)}) \rangle \}$  where:

- $P(I_{ij}^{(k)})$  — price of unit of bandwidth set by ISP  $k$  on the segment  $i,j$
- $B(I_{ij}^{(k)})$  — bandwidth units available from ISP  $k$  on the segment  $i,j$

The key dimensions of connectivity in MINT are *capacity granularity* and *lease time frame*. A transit ISP could also advertise multiple price levels for different lease times (*i.e.*, discounted rate for longer lease periods.) Similarly, transit providers can announce a *minimum reservable bandwidth* parameter, which restricts the minimum amount of bandwidth peers can lease. In a more general approach the announcements and corresponding reservations need not to refer to only bandwidth guarantees. ISPs could also trade loss rates, 95th percentile charges and expected capacity reservations [11].

The mediator maintains a database of all active offers, and uses them to match with available bids when possible. On the other side of the market, source networks are seeking to get low priced paths that match given constraints. A request consists of a source exchange  $s$ , a destination exchange  $d$ , the willingness to pay, and the bandwidth desired. When an incoming bid arrives, the mediator first checks whether any path can be constructed from the existing offers to match the bid; if not, the bid is rejected. If a feasible path exists (*i.e.*, a path with sufficient capacity to match the bid), the mediator constructs the *lowest price feasible path* between  $s$  and  $d$ .

### 4. Protocols

This section details the design goals and the two main components of MINT: the control plane, which is designed in a two-level hierarchy to enable scaling; and the data plane, which establishes end-to-end paths using distributed path setup and tunneling. We explain how each component requires only minimal changes to existing protocols.

#### 4.1 Design Goals

MINT must operate a market for end-to-end connectivity at Internet scale. To do so, it must satisfy the following design requirements:

- **Control-plane scalability.** MINT’s transit-state exchange system must operate at Internet scale. Each independently operated network may contain thousands of routers and links that continually fluctuate, which makes maintaining global link-level state both infeasible and redundant.
- **Data-plane scalability.** MINT uses tunnels to forward traffic along path segments within each independently operated network, “stitching together” their tunnels across domains. These tunnels require additional state for each ingress-egress path. MINT must manage this state and also ensure that paths re-converge quickly when failures occur.
- **Topology privacy.** For competitive or security reasons, ISPs are typically unwilling to disclose any in-

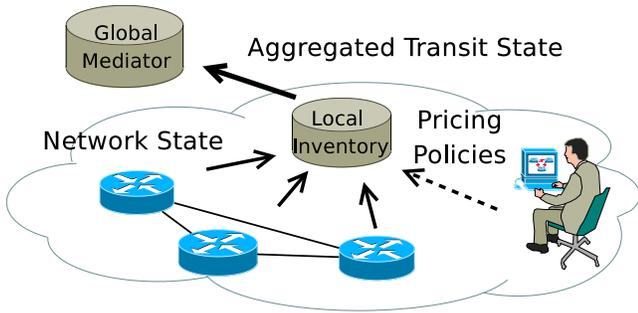


Figure 3: Two-level inventory topology maintenance system.

formation about network topology. MINT must provide abstractions that preserve ISPs’ privacy while still facilitating the sale of path segments.

These design requirements lead to two primary design decisions: (1) a *two-level control-plane hierarchy* (Section 4.2) and (2) *distributed path setup protocols* (Section 4.3). The two-level control hierarchy provides control plane scalability and topology privacy. Distributed path setup should offer data plane scalability and fast convergence.

## 4.2 Control Plane: Two-Level Hierarchy

To scalably provide control-plane information while preserving privacy, MINT uses a two-level hierarchy to disseminate transit information. Within a local network, operators set policies that control the information that is exported to the mediator’s global transit state database. The transit state database registers path segments advertised by each network; this information allows the mediator to perform the exchange-level path computation to satisfy demand for end-to-end paths.

As shown in Figure 3, each network runs a system that performs three main functions: (1) a monitor that tracks the availability and load of local network elements, (2) a management interface by which operators can price segments and set export policies, and (3) an interface to the mediator’s topology service.

**Monitoring connectivity and managing policies.** The monitor can use existing Interior Gateway Protocols (IGPs) to discover information about network availability and allocated bandwidth [19]. The management interface allows operators to react to market fluctuations and set the prices for path segments; we envision that this interface can be implemented using protocols such as SOAP and XML. Similarly, the interface between the local monitor and a global topology service could be implemented using existing publish-subscribe protocols.

**Maintaining global topology information.** Participating domains report the *transit state* for each path segment to a system that maintains global topology. Each transit state is expressed as a set of the ingress and egress exchange points with the properties for the corresponding path segment. Each ISP decides independently which segments to advertise, as well as the capacity and price for each segment. In addition to price and bandwidth, as described in Section 3,

the advertised transit states can contain more detailed information, such as the minimum reservable bandwidth (*i.e.* to enable grooming) and minimum lease time (*i.e.* to offer discount rates for long term users).

**Lazy recovery from failures.** The network state in the global topology database may not always accurately reflect the current state of the underlying physical network. This separation allows MINT to scale while maintaining more network state than today’s routing protocols do. The global transit state database does not automatically recompute interdomain paths when links fail; it only does so when a source requests a new end-to-end path. This lazy recovery approach allows the database to scale, without sacrificing responsiveness. As described in Section 4.3, path failures can be detected by source networks, but only trigger path recomputation as necessary.

In our future research, we will explore two possible directions for computing end-to-end paths. On one hand, the computation of the exchange-level path could be centralized, utilizing powerful server resources (edge networks could be charged for this service). On the other hand, the global topology database could be decentralized, perhaps using flooding in a similar fashion as link-state routing protocols. This approach can scale because forwarding and routing are decoupled, as described in the following section.

## 4.3 Data Plane: Distributed Path Setup

MINT must forward traffic along explicitly reserved end-to-end paths that have been matched to buyers by the mediator. This paradigm represents a departure from conventional IP forwarding, which forwards packets based only on destination. We describe these functions below.

**Forwarding traffic along path segments.** Simple flow-based forwarding schemes needed for end-to-end paths in MINT can be achieved with either hop-by-hop IP-in-IP tunneling or label switching (*e.g.*, MPLS [12]). Such switching for the interdomain scenario is available from many vendors and comes by default with most high-end routing equipment; for example, the cross-domain tunnels can be achieved using “stitching” mechanisms [8]. Today, such technology is primarily used in a large networks that comprise several separate autonomous systems; MINT enables use of this technology to the Internet at large by offering transparent segment trading platform. The scalability of the path stitching depends on the underlying technology in use. Today’s router’s dataplanes, in case of MPLS technology, can support millions of paths. Control plane is shown to scale to thousands of new paths per second [29].

**Establishing tunneling state.** Tunneling requires explicit path establishment at each node along the route, which entails two challenges. First, node-level path computation and reservation consumes scarce processing resources. Second, tunnels can increase the data plane state needed to maintain to implement reservation forwarding. To solve these problems, recent proposals such as the Path Computation Element (PCE) [16] suggest using a dedicated control element local to a given provider to perform such task. This approach reduces the overhead of path computation and setup on indi-

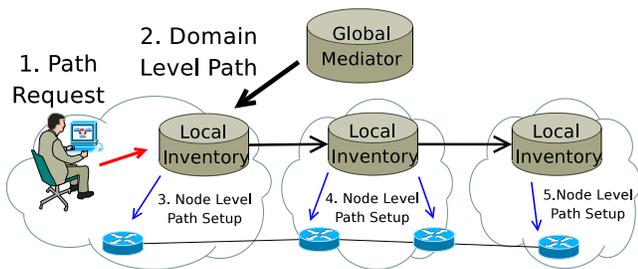


Figure 4: Path computation and setup in MINT. A global mediator computes an end-to-end path that satisfied the user request; local inventories compute and coordinate the setup of node-level path within each ISP

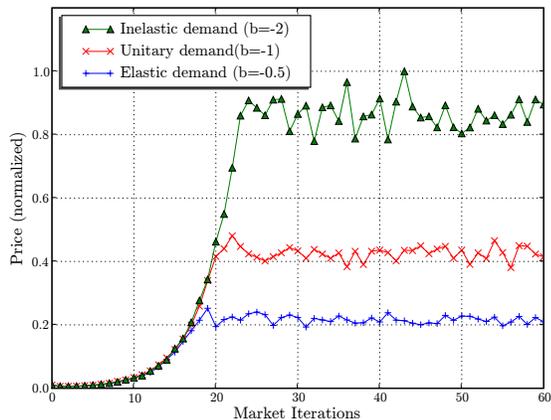


Figure 5: Price convergence in an open bandwidth market. Different plots show price convergence corresponding to different demand functions. Parameter 'b' indicates the slope of the linear demand curve.

vidual routers.

In general, tunneling mechanisms scale well: typical mid-range routers can support tens of thousands of tunnels [26], and these routers can typically scale to millions of tunnels, with setup rates of up to 400,000 tunnels per second [29]. Still, each network's path computation element can also load balance requests across the available resources: if multiple local nodes can satisfy a path request, the element can select nodes that are less loaded.

## 5. Preliminary Evaluation

In this section, we present preliminary evaluation results that study various aspects of MINT's economic viability. We aim to evaluate whether the market in MINT converges. Specifically, we study whether (and how quickly) prices stabilize. We describe our model for an exchange-level topology, traffic demands, and a mediation algorithm.

Requests for bandwidth, or *bids*, between any two exchanges arrive at discrete time intervals according to a Poisson arrival process. Each bid has associated requested bandwidth and a maximum price it is prepared to pay for the path. The maximum price is selected at random from a predetermined, downward sloped demand curve. In experiment we model three types of demand: *elastic*, *unitary* and *inelastic*. In the inelastic case more users are prepared to pay high

price to get the service while in the elastic case users are easy to switch to best-effort traffic instead of paying a premium; unitary demand lies in the middle of the former two. A mediator maintains the database of advertised segments, along with their corresponding prices; these are the *offers*. The bids are processed on a first come, first served basis with a greedy algorithm. For each bid, the mediator runs a constrained shortest path algorithm to find the least expensive path given current offers; if such a path does not exist, or if path price is higher than the price the consumer is prepared to pay, the bid is dropped. Satisfied bids increase the capacity used on each segment of the path.

In a real market, each ISP would set the price for their segments independently, according to their local policies. For our simulation, we develop an algorithm that simulates such behavior. The algorithm works as follows: we start by setting the price of each segment to a small initial value and change that value according to the load. If load is below 80% of the link capacity, we reduce the price, on the other hand, if it is above 80%, we increase the price. More precisely, the price  $P(I_{ij}(t))$  for a segment  $ij$  at a time  $t$  is selected as follows:

$$P(I_{ij}(t)) = \begin{cases} \alpha P(I_{ij}(t-1)), & \text{if } ij \text{ is utilized at } t-1 \\ \alpha^{-1} P(I_{ij}(t-1)), & \text{otherwise} \end{cases}$$

where  $\alpha$  is the rate of the price change. For our experiments, we set  $\alpha = 2$ .

We use data provided by PeeringDB as the basis for a realistic exchange-level topology [4]. PeeringDB contains information about around 200 public Internet exchange points, with details about network participants and their connection speeds. Figure 5 shows the results of MINT simulation using the price selection algorithm described above. Each plot represents a different demand curve and resulting average segment price in a system. All demands models converge to a some average price point within a small number of iterations: inelastic demand produces the highest average prices, while elastic demand produces the lowest prices.

## 6. Deployment Incentives

Beyond the algorithmic and protocol-related deployment concerns, various incentive issues arise. In particular, we must demonstrate that all involved parties benefit from participating in the market.

**Buyers.** Both stub networks and transit providers that procure paths through MINT would benefit from the increased diversity that MINT offers. It is worth exploring, however, whether this is greater diversity than such a network could obtain by other more traditional means. For example, a network might buy upstream connectivity to a very well-connected ISP, and use that ISP to provide reliability benefits for end-to-end paths; Internap provides such a service to downstream customers [17]).

**Sellers.** ISPs with transit capacity must have an incentive to participate in MINT's market. One clear incentive is that providers now have an opportunity to sell capacity that would otherwise have remained dormant. Additionally,

MINT gives long-haul providers more opportunity for differentiation and specialization than today: for example, a long-haul provider might choose to specialize in providing connectivity for certain applications, or between specific geographic regions (rather than providing an expansive backbone).

**Mediator.** The business model for the mediator is similar to that of traditional exchanges and brokerages. Mediator can either collect fixed fees for each service request, or take a percentage fee from the participant revenue.

## 7. Related Work

The connectivity offered by MINT is similar to that provided by various products and services that offer stub networks connectivity via multiple upstream ISPs on short timescales (e.g., Equinix Direct [13]). MINT is also similar to Band-X [2], Arbinet [1] and other PSTN/VoIP exchanges, but provides a more general framework for establishing end-to-end paths for data and addresses issues related to path construction. “Intelligent routing” products [9] perform only path selection on paths that have been established with existing contracts. Unlike most other systems, MINT offers the ability to execute contracts for *end-to-end* paths (rather than with direct upstream ISPs).

Our work has some connection with network pricing [6, 15]. This literature typically follows one of two themes. In some of the work, pricing is used as a mechanism to alleviate congestion inside the network; here, the network is a single entity, and the mechanism aims to control user behavior (see, e.g., the seminal papers of Mackie-Mason and Varian [22] and Odlyzko [24]). A second thread uses congestion pricing as a motivation for the design of distributed congestion control protocols [20, 21, 28]. In these models, the network is also treated as a single entity. By contrast, MINT considers the entire market as service providers and end hosts; competition between service providers plays a key role in determining market outcomes.

## 8. Summary

Today’s markets and protocols for connectivity do not support the wide variety of emerging Internet applications. To provide networks with more flexible access to end-to-end paths, we present MINT: a Market for INternet Transit. MINT allows providers to announce excess capacity along with prices to a mediator, who matches available path segments to bids for end-to-end connectivity. MINT reduces the inefficiencies that exist in today’s connectivity markets, and creates an opportunity for networks to establish end-to-end paths with guarantees and properties otherwise unavailable with conventional routing.

## REFERENCES

[1] Arbinet. <http://www.arbinet.com>.  
 [2] Band-X. <http://www.band-x.com>.  
 [3] Group Suggests an Exchange to Trade Internet Capacity. <http://www.nytimes.com/2008/06/30/technology/30byte.html>.  
 [4] Peering Database. <http://www.peeringdb.com>.

[5] C&W Briefly Drops Peering Agreement with PSINet. [http://www.isp-planet.com/news/2001/psinet\\_cw.html](http://www.isp-planet.com/news/2001/psinet_cw.html), June 2001.  
 [6] M. Afegan and J. Wroclawski. On the Benefits and Feasibility of Incentive Based Routing Infrastructure. In *PINS*, Portland, OR, 2004.  
 [7] Alin Popescu and Todd Underwood. D(3) Peered: Just the Facts Maam. A Technical Review of Level (3)s Depeering of Cogent. In *NANOG 35*, Oct. 2005.  
 [8] A. Ayyangar, K. Kompella, J. Vasseur, and A. Farrel. *Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)*. Internet Engineering Task Force, Feb. 2008. RFC 5150.  
 [9] Cisco Optimized Edge Routing (OER). [http://www.cisco.com/en/US/products/ps6628/products\\_ios\\_protocol\\_option\\_home.html](http://www.cisco.com/en/US/products/ps6628/products_ios_protocol_option_home.html), 2006.  
 [10] Cisco Visual Networking Index. <http://newsroom.cisco.com/visualnetworkingindex/>.  
 [11] D. D. Clark. A model for cost allocation and pricing in the internet. In *Proc. MIT Workshop on Internet Economics*, Mar. 1995.  
 [12] B. Davie and Y. Rekhter. *MPLS: Technology and Applications*. Academic Press, San Diego, CA, 2000.  
 [13] Equinix Direct. [http://www.equinix.com/prod\\_serv/network/ed.htm](http://www.equinix.com/prod_serv/network/ed.htm), 2007.  
 [14] C. Estan, A. Akella, and S. Banerjee. Achieving Good End-to-End Service Using Bill-Pay. In *Proc. 5th ACM Workshop on Hot Topics in Networks (Hotnets-V)*, Irvine, CA, Nov. 2006.  
 [15] C. Estan, A. Akella, and S. Banerjee. A la carte: An Economic Framework for Multi-ISP Service Quality. Technical report, University of Wisconsin, Mar. 2007.  
 [16] A. Farrel, J. Vasseur, and J. Ash. *A Path Computation Element (PCE)-Based Architecture*. Internet Engineering Task Force, Aug. 2006. RFC 4655.  
 [17] Internap. <http://www.internap.com/>, 2006.  
 [18] Invisible Hand Networks. <http://www.invisiblehand.net/>.  
 [19] D. Katz, K. Kompella, and D. Yeung. *Traffic Engineering (TE) Extensions to OSPF Version 2*. Internet Engineering Task Force, Sept. 2003. RFC 3630.  
 [20] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.  
 [21] F. P. Kelly, A. K. Maulloo, and D. K. Tan. Rate control for communication networks: shadow prices, proportional fairness, and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.  
 [22] J. K. MacKie-Mason and H. R. Varian. Pricing the Internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*, pages 269–314. MIT Press, Cambridge, Massachusetts, 1995.  
 [23] Martin A. Brown and Alin Popescu and Earl Zmijewski. Peering Wars: Lessons Learned from the Cogent-Telia De-peering. In *NANOG 43*, June 2008.  
 [24] A. M. Odlyzko. Paris Metro Pricing for the Internet. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 140–147, 1999.  
 [25] S. Parsons et al. Everything you wanted to know about double auctions but were afraid to (bid or) ask. <http://www.sci.brooklyn.cuny.edu/~parsons/projects/mech-design/publications/cda.pdf>, 2006.  
 [26] R. Raszuk. Mpls traffic engineering overview. <http://www.nanog.org/mtg-0002/robert.html>, Feb. 2000. NANOG18.  
 [27] Y. Rekhter, T. Li, and S. Hares. *A Border Gateway Protocol 4 (BGP-4)*. Internet Engineering Task Force, Jan. 2006. RFC 4271.  
 [28] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, Boston, Massachusetts, 2004.  
 [29] H. Wang, R. Karri, M. Veeraraghavan, and T. Li. A hardware-accelerated implementation of the RSVP-TE signaling protocol. In *Proc. IEEE International Conference on Communications*, June 2004.  
 [30] Z.-L. Zhang. Decoupling QoS control from core routers: A novel bandwidth broker architecture for scalable support of guaranteed services. In *Proc. ACM SIGCOMM*, pages 71–83, Stockholm, Sweden, Sept. 2000.