RZ 3294 (# 93340) 11/27/00
Computer Science 11 pages

Research Report

A Market-Based Model of Bandwidth and a New Approach to End-to-End Path Computation with Quality Guarantees

G. Cheliotis

IBM Research
Zurich Research Laboratory
Säumerstrasse 4
8803 Rüschlikon
Switzerland

Research
Almaden · Austin · Beijing · Delhi · Haifa · T.J. Watson · Tokyo · Zurich

A Market-Based Model of Bandwidth and a New Approach to End-to-End Path Computation with Quality Guarantees

G. Cheliotis

IBM Research, Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

Abstract

A market-based end-to-end IP service model with quality-of-service (QoS) is proposed based on a recently introduced commodity contract for bandwidth. A percentage of the total trading volume in the bandwidth market is traded in terms of standard commodity contracts and across multiple public marketplaces. The various roles in this model and their relationships are explained and the notion of a bandwidth contract graph is introduced. The problem of selecting an optimal set of bandwidth contracts for end-to-end connectivity that satisfies a budget constraint and a set of quality requirements is shown to be a link-constrained multi-path-constrained path optimization problem. This includes the multi-path-constrained problem (MPC), whose general formulation is known to be NP-complete, but we show that polynomial-time algorithms are applicable for standardized contracts. Standard contracts thus reduce the computational complexity, facilitating the expansion of bandwidth markets and the development of new, efficient brokerage services for end-to-end connectivity with quality guarantees.

1 Introduction

Computer networks interconnect partly due to state regulation, but also to take advantage of the positive network externalities inherent in computer networks and services. There may be concerns at to whether it is profitable for a network provider to interconnect with another provider under specific circumstances (studied in [BW99, JCT99, Hus99]), but the current main drivers for interconnection, i.e. universal access and universal services, generate especially strong positive feedback [SV99], which guarantees that network and service providers will seek ways to expand the reach of their networks. The question is how to do that, how to provide quality guarantees, what rules and agreements should govern the interconnection relationships and last but not least, how to price bandwidth.

Currently there are three basic ways to expand the reach of a network, either by interconnecting with other networks[1], thus buying the right to send traffic through those networks, or by directly buying capacity on those networks. We will briefly present these options in chronological order of appearance.

Private Interconnection Points. Network interconnection points with bilateral financial settlements or without (the latter also known as "peering"). Fully interconnecting N providers means provisioning $N*(N-1)/2$ connections. In Fig. 1a three network clouds and their private interconnection points are shown.

Public Interconnection Points. A different solution is the creation of an independent peering point, owned and operated by a consortium of providers or a trusted third party. Providers have to connect to that single point. In this case, interconnecting N providers requires N connections to the peering point (see Fig. 1b). This has been the most popular model of interconnection in the public Internet and it is usually associated with settlement-free peering agreements.

What is common to the two approaches above is that the set of participants is often limited to a few carriers, and long-term bilateral or multilateral agreements are the norm. Furthermore, in both cases the issues of pricing and charging for the interconnection benefits have no simple, unique solution [Hus99].

Commodity Markets. Recently a third way to expand a network's reach was introduced, based on the concept of a public market for connectivity, similar in nature to commodity markets (e.g. crude oil or electricity). As shown in Fig. 1c, the view of an interconnection point has changed significantly; all a member of such a market sees (or indeed cares to see) is a list of offers from different providers for particular network destinations.[2] In Fig. 1c, circles represent destinations (i.e. exchange locations) and lines stand for contracts between X and the destination exchange points, termed "bandwidth segment contracts", offered by the providers/networks in braces. The providers themselves might not even be visible at all. In the case where offers are equivalent with respect to some criteria, or benchmarked according to an industry standard, choice can be based on destination and price alone. This is a "public" market open to anyone who has registered and is legally entitled to trade goods, producers,

---

[1] See [Bai97] for an overview of current interconnection architectures.

[2] A destination in this paper is a traffic exchange location, where a buyer's port on a switch/router can be connected to the seller's port for routing traffic after a deal has been struck in the market. These destinations were "hidden" behind the clouds of cases (a) and (b) of Fig. 1.

consumers of bandwidth and speculators alike. Arguments in support of this market-based model include increased bandwidth liquidity, clearly specified and widely used contracts, transparent and competitive pricing and simplified charging (compared to financial settlements in public and private peering points) [Enr99, dPG98].[3]

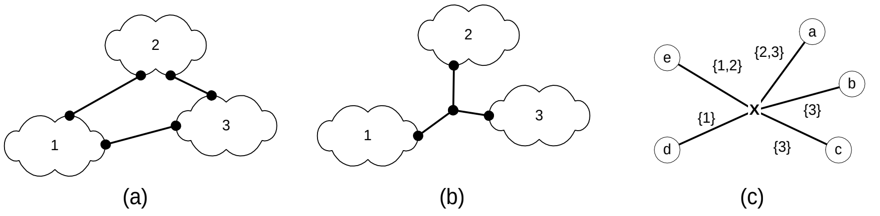

Figure 1: Expanding a network's reach: (a) private peering points, (b) multilateral peering and (c) commodity markets.

This paper concentrates on bandwidth commodity markets and defines an extended version of the above desription for trading bandwidth across multiple exchange points. The extended model is a natural evolution of current efforts to introduce bandwidth trading at a couple of geographical locations and partly draws on the ways established commodity markets operate.

It is a common practice to define industry-wide standard contracts in commodity markets and to encourage traders to trade using these contracts [dPG98]. Standard contracts are chosen such that they are representative of different variants, "grades" of the commodity available in the market. The trading volume of standard contracts for a particular commodity need not be high; their spot prices act as indicators for the pricing of all contracts of this commodity. Non-standard contracts are benchmarked relatively to the standard contracts and priced accordingly. There is no need for a well-defined mapping of quality differentials to price differentials, it is competition and public knowledge of the spot prices which force providers of the commodity to price relatively to the standard contracts.

In the case of bandwidth commodities the natural quality differentiator is quality-of-service (QoS). There is also a natural segmentation of the bandwidth market into customers with different bandwidth size requirements, so the amount of bandwidth is what defines the "grades" of the commodity (e.g. DS-3, OC-12).[4] In this paper it is assumed that some standard contracts with a defined bandwidth and quality level (e.g. "DS-3 corporate", "OC-12 carrier")[5] are traded. We show how their pricing information helps solve otherwise difficult problems.

The remainder of this paper is oganized as follows: Section 2 describes in detail the layered model of the bandwidth market, Section 3 defines the general path computation problem,

---

[3] Throughout this paper it is assumed that some price setting mechanism is in place at marketplaces, matching bids and offers. This could be for example a double auction. See [San99] for an example of a bandwidth auction.

[4] Although acronyms such as DS-3 actually refer to different transport technologies, they are often used in telecommunications to denote a bandwidth size. For example, in [Enr99] a DS-3 equivalent (44.736 Mbps) is defined.

[5] The implicit assumption here is that bandwidth and quality are coupled in the standard contracts, i.e. a OC-12 standard contract comes along with a specified quality level. This is to simplify the presentation of some concepts later on. Extending to two quality levels (e.g. one for voice and one for data traffic) per bandwidth level is simple and can be dealt with by simply running the algorithms in Section 5 twice, or by requiring that a customer states his application.

and Section 4 describes the problem in more detail. Section 5 proposes two algorithms and illustrates their use with a brief example. The paper concludes with an evaluation of the findings and identifies some issues for future research.

2 A Layered Model of Connectivity Trading

In the new bandwidth trading model we define the roles of a (connectivity) provider, a broker, a customer, an exchange and a market data provider. An organization could assume more than one role, e.g. a provider is often a customer of other providers and could also function as a broker.

- A *Bandwidth eXchange* is a local marketplace for suppliers (providers) and buyers of bandwidth (customers) and a clearinghouse. The latter, termed “pooling point” in [Enr99], is a switch/router or a small network thereof for the actual bandwidth delivery upon settlement. The marketplace and clearinghouse could be operated by different parties specialized in trading and networking, respectively. In this paper we use the terms “bandwidth exchange” and “exchange” interchangeably. An “exchange location” is the geographical location of the clearinghouse (i.e. where network connections are actually made). It is assumed that an exchange trades well-defined bandwidth contracts. Each contract refers to a *bandwidth segment* between two exchange locations. A bandwidth segment is an abstraction of one or more providers’ networks offering connectivity from one exchange location to another.

- A *Provider* offers connectivity from one exchange location to another in the form of bandwidth *segment contracts*. Providers have to be connected to the clearinghouses of both exchanges to make offers on the segment defined by the two exchange locations.

- A *Customer* generates requests for end-to-end paths. We are not concerned here with traffic origination and termination beyond the exchange locations. In this context an end-to-end path involves only exchange locations and their clearinghouses.

- A *Broker* processes customer requests and computes end-to-end paths meeting the customer’s requirements at the lowest possible cost. The broker is an important figure in this text because of his central role in the configuration of paths. Prices of contracts are cost elements for a broker and his task is to minimize the total cost on paths. Brokers must be subscribed to at least one MDP (see definition below) to get up-to-date data for all traded segments.

- A *Market Data Provider (MDP)* is connected to all the bandwidth exchanges and gets all current market data from them. He then offers this global market information along with some value-add (e.g. market-wide statistics) as a service to subscribed brokers.

A distinction is made in this paper between a network and a trading layer. This distinction is useful because one can choose to approach the model from one or the other layer and will either way get a coherent view of interconnection, albeit from a different angle.

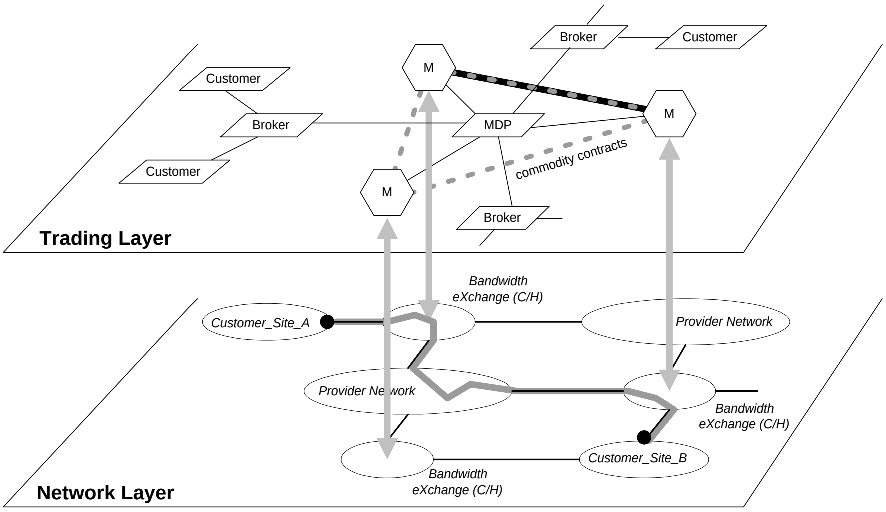

Figure 2: The market-based bandwidth model. The bold black line represents a segment contract.

Network Layer. The network layer looks very much the same as some current inter-domain network topologies[6] where multiple carriers interconnect at certain public or private peering points, with the exception of the change in terminology: now the peering point is a bandwidth exchange – to be more precise, it is a bandwidth clearinghouse (C/H). A connectivity provider's natural field of operation is the network layer. Figure 2 shows an inter-domain topology of two local customer networks, two provider networks and three clearinghouses of bandwidth exchanges. Straight lines represent inter-domain network links. The bold gray line between the two customer sites represents a configured inter- and intra-domain path for the connection of the two customer networks. At each exchange location traffic is switched according to contract relationships[7] depicted in the trading layer. A provider is free to choose any resource reservation scheme internally to meet the bandwidth and quality requirements of the constracts he sells.

Trading Layer. This layer looks quite different and the main reason for this is that provider networks are "hidden" behind contracts for bandwidth segments traded at exchanges. Each contract defines, among other things, capacity (or a capacity equivalent) and a set of benchmark QoS parameters (see [Enr99] for a detailed example). Brokers, MDPs and customers are mainly concerned with this layer in their daily market transactions. Figure 2 shows how clearinghouses (C/H) map to local marketplaces (M) and the path in the network layer (the inter-exchange part of the path) maps to a segment contract depicted as a bold black line in the trading layer. This contract is a link in the *contract graph* whose nodes are the local marketplaces of bandwidth exchanges.[8] These concepts will be explained in more detail later on. In Figure 2 one can also see how a trading network of brokers, MDPs and customers is created in this layer.

---

[6] A domain here is a network or set of networks under the same administration.

[7] A proposal for achieving this using layer-3 switching is included in [Enr99].

[8] A link between A and B in the contract graph represents the fact that the segment defined by the exchange locations A and B is being traded (at the respective marketplaces).

3 Basic Definitions

We are now more formally defining some basic concepts which are useful for formulating the path computation problem, the solution of which is actually the broker's main task.

3.1 Definitions

A bandwidth segment s is a pair of exchange locations (a, b) corresponding to neighboring vertices of the contract graph. We define a contract $\mathcal{C}$ for segment s as follows: $\mathcal{C}_s = (\beta, G(\mathbf{q}), p)$, where β is the bandwidth offered, $G(\mathbf{q})$ is a set of guarantees (i.e. constraints) on a QoS vector $\mathbf{q}$ and p is the price of the contract. The set of all traded segments can be represented as a *contract graph* $\mathcal{G}(V, E)$, where vertices represent exchange locations and edges represent traded bandwidth segments. Contract graphs are in general directed because a contract may specify different guarantees for the two directions between too nodes (although this is not the case in this paper). A path π in the contract graph is a sequence of exchange locations $\pi =< a_1, a_2, ..., a_n >, \quad a_i \in \mathcal{G}$.[9]

3.2 The Brokerage Problem

A customer's request is expressed in terms of a pair of exchange locations (a, z), a bandwidth value β', and either an upper price bound p' or a set of end-to-end QoS requirements $R(\mathbf{q}(\pi))$, or both. The problem the broker is facing is to compute a path π in the contract graph G, such that the first and last elements of π are a and z respectively, the bandwidth β' requested can be provided by π, path gurantees $G(\mathbf{q}(\pi))$ satisfy $R(\mathbf{q}(\pi))$ and $p(\pi) \leq p'$. The details of these conditions are discussed in the next section.

The above general description bears some resemblance to QoS routing problems [CN98b]. QoS routing in general is handled at the intra- and inter-domain levels, where a domain is defined as a network or a set of networks under a common administration, also called "administrative domain". Intra-domain QoS support is the responsibility of the administrative entity (provider) in charge of that domain and depends heavily on particular technological characteristics of the networks in that domain. Inter-domain QoS routing, on the other hand, is more about bridging the gaps between the different technologies and quality specifications in the domains, to enable seamless connectivity with QoS over multiple domains [KKB$^+$98]. Our market-based model leads to a reformulation of the inter-domain QoS routing problem, mainly due to the fact that we introduce the bandwidth segment and contract graph abstractions, as described in Sections 1 and 2. This formulation is different from inter-domain QoS routing in general and it leads to a computationaly efficient solution under certain reasonable assumptions about contracts.

4 Path Computation

We continue with a closer examination of the path computation problem. Our graph consists of segments and we assume there are available contracts traded for each segment. It has been suggested in [Enr99] and [dPG98] to introduce standard bandwidth contracts. As it is

---

[9] What does the contract graph look like? In early stages of bandwidth trading it is bound to be limited in size, but when the market matures and expands, the graph will grow and the connectivity degree will increase.

generally the case that bandwidth is offered in standard increments (e.g. DS-3, OC-12, etc) and quality differentiation is often a matter of a couple of QoS classes, it is fair to assume that a couple of standard bandwidth contracts can be defined and will be traded (see also previous discussion on commodity markets). This, as we will see, is helpful in determining the properties of a path in the contract graph.

4.1 Segment and Path Properties

In Section 3 a segment contract is defined as the triad $(\beta, G(\mathbf{q}), p)$. Let us take a set of commonly used parameters, often referred to in service-level agreements, as our QoS vector $\mathbf{q} =< d, \Delta d, L, R >^T$, where d is delay, Δd is delay variation (jitter), L is the loss ratio (traffic lost over traffic sent) and R stands for reliability (1 minus probability of failure). A QoS vector can characterize a segment or a whole path. When computing the elements of $\mathbf{q}$ for a path we should bear in mind that delay and jitter are additive metrics, reliability is multiplicative [WC96, CN98b] and loss ratio is quasi-multiplicative, as shown below. Bandwidth β is a concave metric [WC96] and price p is additive. Our previous problem formulation in Section 3.2 is a *link-constrained (bandwidth) multi-path-constrained (quality, price) path computation* problem which includes the multi-path-constrained routing problem (MPC) [CN98b]; the latter proved to be NP-complete [WC96, GJ79].

The guarantees of a segment contract $G(\mathbf{q})$ can be expressed in terms of the following relations: $d \leq M_d$, $\Delta d \leq M_{\Delta d}$, $L \leq M_L$ and $R = F_R$, where M_x denotes an upper bound of property x and F_x is the value of an exact guarantee. For simplicity we assume that the bandwidth property has an exact value guaranteed (i.e. in our notation $\beta = F_\beta$) which reads as: maximum guaranteed bandwidth is F_β. Note that such a bandwidth guarantee can be provided on any network supporting properly engineered resource reservations. It is the responsibility of the connectivity provider to employ a reservation scheme that is suitable for his network (per-flow reservations, relative differentation of flow aggregates, etc). This is beyond the scope of this paper (see [FH98] for a discussion of QoS mechanisms).

A path's properties are computed from a single segment's values using the following relations:

$$\begin{aligned}
\beta(\pi) &= min_{i \in \pi}\{F_{\beta_i}\} && (1)\\
p(\pi) &= \sum_{i \in \pi} p_i && (2)\\
d(\pi) &\leq \sum_{i \in \pi} M_{di} && (3)\\
\Delta d(\pi) &\leq \sum_{i \in \pi} M_{\Delta di} && (4)\\
L(\pi) &\leq 1 - \prod_{i \in \pi}(1 - M_{Li}) && (5)\\
R(\pi) &= \prod_{i \in \pi} F_{Ri} && (6)
\end{aligned}$$

Relations (1) to (6) can be used to compute the properties of a *path contract.* Now let us assume all segment contracts are identical — some are still more expensive than others —, i.e.

our path consists of one standard type of contracts only.[10] Then relations (1) to (6) become:

$$\begin{aligned} \beta(\pi) &= F_\beta & (7) \\ p(\pi) &= \sum_{i\in\pi} p_i & (8) \\ d(\pi) &\leq nM_d & (9) \\ \Delta d(\pi) &\leq nM_{\Delta d} & (10) \\ L(\pi) &\leq 1-(1-M_L)^n & (11) \\ R(\pi) &= {F_R}^n & (12) \end{aligned}$$

where n is the number of segment contracts in the path. If instead of the loss ratio L we use $S = 1 - L$, the success ratio (traffic delivered over traffic sent between two points on a network), inequality (11) is equivalent to the following:

$$S(\pi) \geq {M_S}^n \qquad (13)$$

The above relations show that bandwidth becomes an invariant, the price formula has not changed, and all other properties of the path contract are clearly getting worse the greater the number of segment contracts n in the path. Furthermore, all quality parameters are expressed as simple linear or exponential functions of n. Note that because the bandwidth value remains fixed, it becomes irrelevant in this case. It is only relevant in determining the "grade" of standard contracts to be considered.

We can represent a customer's end-to-end quality requirements $R(\mathbf{q}(\pi))$ with the following bounds on values of the path's QoS vector : $M_{d(\pi)}$, $M_{\Delta d(\pi)}$, $M_{L(\pi)}$, $m_{R(\pi)}$, where m_x is a lower bound on property x. Relations (9) to (13) together with the inequalities in $R(\mathbf{q}(\pi))$ yield

$$\begin{aligned} nM_d &\leq M_{d(\pi)} & (14) \\ nM_{\Delta d} &\leq M_{\Delta d(\pi)} & (15) \\ {M_S}^n &\geq 1 - M_{L(\pi)} & (16) \\ {F_R}^n &\geq m_{R(\pi)} & (17) \end{aligned}$$

Let $\mathcal{S}$ be the solution space of the system of inequalities (14) to (17). By choosing $\hat{n} = \max(\mathcal{S})$, customer quality requirements can be met by a path consisting of $\hat{n}$ segment contracts (or less). Thus we have shown that the entire quality vector can be reduced to an upper bound for the number of contracts in a path, assuming identical contracts. The number of contracts n can be thought of as the "hop-count" in the contract graph.

5 Solutions

The above discussion of path properties shows that under the assumption of identical segment contracts being traded on all segments (or on selected segments indicative of the pricing in an area) we can reduce the general brokerage problem of Section 3.2 to the following least-cost

---

[10] See previous discussion on commoditization of bandwidth.

hop-constrained path optimization problem (LCHC): *find a path in the contract graph that minimizes cost subject to a hop-count constraint.* If the resulting path's cost is lower than the customer's price requirement, then there is a solution (broker operating at zero profit). Alternatively, we could try to minimize hop-count under a cost constraint. In part it depends on which constraints the customer has given.

5.1 Computing the Optimal Path

Both of these problems are path-constrained path optimization (PCPO) problems with two additive metrics, which are NP-complete in general [WC96, GJ79]. They can both be solved using the approach in [CN98a], where one of the additive metrics is discretized and an $O(x^2V^2)$ algorithm is given for minimizing the other metric subject to the constraint on the discrete one. In that algorithm, x is an approximation parameter which at higher values increases the success of the algorithm at an increased complexity cost. Space complexity is $O(xV)$.

Chen and Nahrstedt suggest in [CN98a] the use of an approximation technique, which, in fact, is not necessary for our LCHC problem because the hop-count is constant at every edge of the graph and therefore does not need to be discretized. Without the approximation step, LCHC can be solved in $O(\hat{n}^2V^2)$ using the extended Dijkstra shortest path algorithm[11] in [CN98a], which is guaranteed to find a solution, if there is one. Space complexity in this case is $O(\hat{n}V)$.

5.2 Computing a Hint

Instead of aiming for the optimal solution, a broker can also give the customer a quick hint for a requested path, requiring less information from the customer's side and using a more efficient algorithm than above. Assume a customer asks for a price for bandwidth β'' from a to z (e.g. "What's the price for a DS-3 equivalent from London to New York?").

Assuming there is an upper bound Φ for prices of all traded segments (this could be the currently highest price plus a margin), we can define the mixed metric $w(u,v) = p(u,v) + \Phi$, which is additive. The term Φ represents the hop-count, scaled to the level of prices in the mixed metric. Using this as the only distance metric, we can look for the shortest path in the contract graph using a plain Dijkstra algorithm, with time complexity $O(|E|+|V|\log|V|)$ when implemented using a Fibonacci heap as a priority queue [Jun94]. This bound is better than the complexity in the previous section, especially in the case of graphs with a low connectivity degree. The path π^* computed (if any) represents an almost equal tradeoff between price and hop-count (assuming low price dispersion), but with a bias towards reducing hop-count. Path π^* is the cheapest among all possible paths from a to z with the same hop-count n^*, and there is no cheaper path with a lower hop-count. Going back to the standard contracts and choosing the β''-grade contract, we can compute the quality guarantees $G(\mathbf{q}(\pi^*))$ of the path from relations (7) to (13) (where $n = n^*$). Then the "hint" to the customer is: "The cheapest path from a to z with β'' bandwidth and quality guarantees $G(\mathbf{q}(\pi^*))$ costs $p(\pi^*)$". This gives a fair indication of the path's current price and if the hint itself is satisfactory to the customer he can buy the contracts in the computed path. If not, he must supply his quality requirements and/or upper price bound, and the algorithm of the previous subsection will attempt to find a qualifying path at a higher complexity cost.

---

[11] The extended Bellman-Ford algorithm can also be used.

5.3 Short Example

Consider the contract graph of Figure 3. Suppose a customer asks for a "hint" for a DS-3 bandwidth equivalent from A to F. His broker, using the mixed metric $w(u,v) = p(u,v) + 9$ with Dijkstra's shortest-path algorithm, finds the path $\pi^* =< A, C, D, F >$, which is the cheapest path with hop-count $n^* = 3$. Then the broker looks up the quality guarantees $G(\mathbf{q})$ of the standard DS-3-grade segment contract to compute $G(\mathbf{q}(\pi^*))$.

Suppose the customer is not satisfied with the quality guarantees of the path, so he makes a more detailed request, specifying his quality requirements $G(\mathbf{q}(\pi))$ for the path. The broker can now combine these requirements with the DS-3-grade segment contract guarantees, which leads to the system of inequalities (14) to (17). Solving this system results in a maximum hop-count $\hat{n}$, which satisfies all quality requirements. Suppose $\hat{n} = 2$. Then the broker has to run the extended Dijkstra algorithm to find the cheapest path with hop-count not greater than 2. The result is $\hat{\pi} =< A, E, F >$. Now the answer to the customer is that there is a path that satisfies his quality requirements. It consists of two standard DS-3 segment contracts and its price is 17.

If the price is within the customer's budget he buys the two standard DS-3 segment contracts. If not, he knows he may find cheaper offers for DS-3 from A to Z traded in non-standard contracts. However, because pricing of these offers is expected to be relative to the spot prices of the "benchmark" standard DS-3 contract, these will be lower-quality offers and maybe he will have to accept a quality compromise after all.

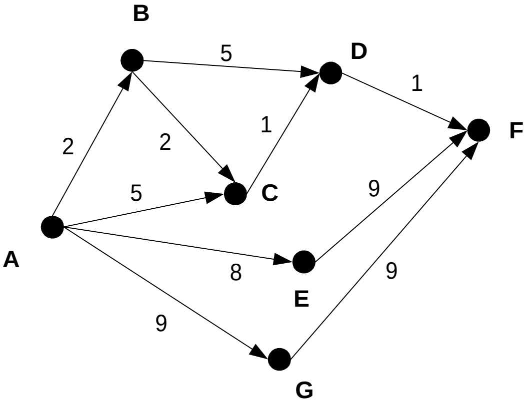

Figure 3: An example of a contract graph.

6 Conclusions and Future Work

The introduction of the market model for bandwidth trading and the layered views assist in the understanding of the interplay between markets and networks and how the economics of service quality will influence the future of IP networks. Interesting issues arise when the leap from individual bandwidth segment contracts to end-to-end path contracts is made, leading to a reformulation of the inter-domain QoS routing problem. It was argued that, using the new formulation, the computation of end-to-end paths with quality guarantees can be offered as a new form of brokerage service to customers.

Furthermore it was shown that the standardization of bandwidth contracts leads to a lower complexity for quality-constrained path computation. This means that the definition of market-wide standard commodity contracts will enable the development of new brokerage services in the bandwidth market and forward the adoption of internet service quality mechanisms. The standard commodity contracts need constitute only a small percentage of the total bandwidth trading volume and behave as market indicators. Standardization in this case does not imply standards imposed by a regulatory body or other authority. On the contrary, we expect proposals such as in [Enr99] to be made to the market and the de facto commodity standards will be those that market participants will trade the most. This form of market-driven standardization is shown to be beneficial for the delivery of IP service quality, even when it is limited to a small perecentage of the world's network capacity.

As bandwidth itself is turning into a tradable good, the boundaries between networks and electronic markets are no longer so distinct. Du Pre Gauntt states in [dPG98] that the "network is the market", but using the bandwidth segment abstraction we could also claim that the market is the network. This can be interpreted in the following manner: assuming a public, mature and liquid bandwidth market with real-time trading, an organization with a privately owned network and a couple of connections to exchange locations of our model can regard the whole set of traded bandwidth segments as a virtual extension of its own network because it can choose to expand at any time by buying contracts on those segments and prices for all segments are known beforehand. In the same way a privately owned network link has some known costs associated with it (one-time setup and variable maintenance costs), a segment contract has a price that is also known. They both serve the same purpose of connectivity with the difference that a segment contract can be bought anytime at will and at the market price, if and when that extra amount of bandwidth is needed (no direct sunk costs, no costs beyond the contract duration). When a customer contacts a provider directly for a connectivity service, the provider can consider the union of his private network and the spot market as the complete set of alternatives for making an offer, evaluating the pros and cons of using capacity on the private network versus buying a segment or a path on the market.

Topics for future research work include algorithms for path computation taking advantage of particular market characteristics or contract graph topologies, dealing with the uncertainty of price information, refining the model, dealing with the effects of speculation in a public market and developing scenarios of interactions between parties involved in the trading process. Identifying pricing or quality relations between commoditized and non-commoditized bandwidth products may also prove to be an interesting area of research.

7 Acknowledgements

The author thanks Beat Liver for his continuous support and Gabriel Dermler for valuable discussions.

References

[Bai97] J.P. Bailey. The economics of internet interconnection agreements. In L.W. McKnight and J.P. Bailey, editors, *Internet Economics*, pages 155–168. MIT Press, 1997.

[BW99] P. Baake and T. Wichmann. On the economics of internet peering. *Netnomics 1*, pages 89–105, 1999.

[CN98a] S. Chen and K. Nahrstedt. On finding multi-constrained paths. In *IEEE International Conference on Communications (ICC 98)*, 1998.

[CN98b] S. Chen and K. Nahrstedt. An overview of quality of service routing for next-generation high-speed networks: Problems and solutions. *IEEE Network*, pages 64–79, November/December 1998.

[dPG98] J. du Pre Gauntt. The network is the market: Financing internet bandwidth. In *Inet 98*, 1998.

[Enr99] Enron Communications. *Bandwidth Commodity: Market Starter Kit (V 1.2 6 99)*, 1999.

[FH98] P. Ferguson and G. Huston. Quality of service on the internet: Fact, fiction or compromise? In *Inet 98*, 1998.

[GJ79] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Co, San Francisco, 1979.

[Hus99] G. Huston. Interconnection, peering and settlements (on-line article). http://www.telstra.net/gih/peerdocs/peer.html, January 1999.

[JCT99] P. Rey J. Crémer and J. Tirole. Connectivity in the commercial internet. Technical report, Institut d' Economie Industrielle, Toulouse, May 1999.

[Jun94] D. Jungnickel. *Graphen, Netzwerke und Algorithmen (3. Auflage).* Bibliographisches Institut und F.A. Brockhaus AG, Mannheim, 1994.

[KKB+98] D. Karali, F. Karayannis, K. Berdekas, J. Reilly, and D. Romano-Critchley. Qos-based multi-domain routing in public broadband networks. In *5th International Conference on Intelligence in Services and Networks*, 1998.

[San99] T. Sandholm. emediator: A next generation electronic commerce server. In *AAAI '99 Workshop: AI in Electronic Commerce*, 1999.

[SV99] C. Shapiro and H.R. Varian. *Information Rules: A Stategic Guide to the Network Economy.* HBS Press, Boston MA, 1999.

[WC96] Z. Wang and J. Crowcroft. Quality-of-service routing for supporting multimedia applications. *IEEE Journal on Selected Areas In Communications*, pages 1228–1234, September 1996.