

RZ 3244 (# 93290) 12/06/00
Computer Science/Mathematics 12 pages

Research Report

Bandwidth Trading in the Real World: Findings and Implications for Commodities Brokerage

Giorgos Cheliotis

IBM Research
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties).

Bandwidth Trading in the Real World: Findings and Implications for Commodities Brokerage

Giorgos Cheliotis

IBM Research, Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

Abstract

In [Che] a trading model for bandwidth commodities is introduced and the computational complexity of brokerage in a spot market is analyzed. In this model the bandwidth brokerage problem is defined as the packaging of multiple contracts traded in local bandwidth marketplaces to produce new composite contracts. It is argued in the same paper that brokerage of bandwidth commodities is a different approach to the inter-domain QoS routing problem. Based on the insights of the previous work and an analysis of a limited set of recent real-world bandwidth trading data, we proceed to a discussion of the effects of choosing one or the other set of data for the computation of composite contracts. Some general observations are made about the state of the bandwidth commodity market today. It is shown with examples that the choice of input data for brokerage decision support tools is affecting the management of various identified risk factors. The degree to which the broker is exposed to risk when acting as a reseller depends in turn on the choice of paths in a “contract graph” representing the state of the bandwidth market.

1 Introduction

In recent years advances in switching/routing and network management have lifted some of the technological barriers in provisioning of bandwidth on-demand. Meanwhile, according to [CMMG97], the unprecedented abundance of bandwidth supply witnessed in the late nineties will be long-lasting. This trend will urge suppliers and consumers to engage more actively in trading excess capacity, leading to the formation of bandwidth commodity markets. Indeed, several efforts to promote bandwidth trading and the definition of standard contracts are underway today, as in [Enr99].

In [Che] the concept of a bandwidth broker is presented, implementing the basic brokerage functions of searching, providing trust, aggregation and negotiation, as described in [SB99], as well as functions specific to bandwidth markets. The latter include trading multiple contracts in a spot market and even offering new composite contracts to customers, thus creating a secondary market for bandwidth [CL00]. A contract graph is introduced in the same paper to represent the state of the bandwidth market. It is assumed that such a graph consists of nodes representing local bandwidth exchanges and links corresponding to the traded segments, i.e. a link from A to B in the contract graph represents the fact that connectivity between the A and B exchange locations is being traded. Each link was assigned a price tag and the brokerage problem was defined as follows:

Given a customer request for connectivity between two locations with a set of quality constraints and a price constraint, compute a path between these two locations in the contract graph, such that the bandwidth requested can be supported by the path and quality and price constraints are satisfied.

The path found, if any, determines the properties of a *path contract* which the broker may choose to offer to a requesting customer, or may publish on the secondary market. Trading individual segments is a special case of path brokerage, segments being paths of unit length.

It was argued in [Che] that the above problem description is a link-constrained (bandwidth) multi-path-constrained (delay, reliability, etc.) path computation problem, known in the QoS routing literature [CN98]. The general formulation of the problem is NP-complete [WC96, GJ79]. In the case, however, where a set of standard commodity contracts are traded for all segments of the bandwidth market, the brokerage problem was shown to be reducible to a least-cost hop-count-constrained path computation problem. The latter is solved in polynomial time, by applying an extended shortest path algorithm. Thus contract standardization enables automated, scalable bandwidth trading.

This paper is organized in the following manner: Section 2 briefly illustrates the existence of risk in bandwidth brokerage. Section 3 describes the structure and properties of contract graphs. Section 4 lists several types of market data which are useful in bandwidth trading and Section 5 proposes some ways of using this data in the contract graph structure. A large part of the paper is devoted to Section 6, on the collection and evaluation of real-world trading data.

2 Brokerage Risk

The choice of a path contract and therefore the broker's risk associated with trading it, depend on the data contained in the contract graph when computing the optimal path. Risk arises from the fact that market prices and liquidity vary, so a path contract's price and availability is also subject to change. Three risk factors have been identified:

- *Price (Market) Risk* is the risk associated with the change in value of a trader's positions caused by changes in market supply and demand.
- *Liquidity Risk* is stemming from the lack of marketability of a contract at the current market price and is usually reflected in a wide bid-ask spread and large price movements in response to any attempt to buy or sell [IFC].
- *Quality Risk* is an effect of the lack of fully standardized contracts and is reflected in wide ask spreads. In a market where quality is important yet not always well-defined, measurable or comparable, sellers may ask for widely varied prices, on the grounds of brand reputation or simply taking advantage of incomplete information on the buyers' side.

All three of the risk factors are interrelated and therefore not always distinct in this paper. For example, liquidity risk is affecting the price at which the liquidation of a large position could take place, while contract standardization would eliminate quality risk *and* reduce liquidity risk.

In one possible scenario, assume a broker is publishing a forward contract for a path to the market and postpones the purchase of the underlying segments until an order/confirmation is received from a customer, or even until shortly before the time of delivery (just-in-time resource allocation). In these cases, if buying the associated segments turns out to be more expensive than the agreed contract price, the broker will lose money. If, on the other hand, a broker attempts to minimize this risk by over-estimating a path's price in his offer, or by means of early resource allocation, the risk of losing customers to competing brokers will force the broker to "correct" his estimates and seek an optimal allocation strategy. We will examine different ways of computing a path, depending on the set of market data that is available to the broker and show with examples how a broker's exposure to risk is affected by his choice of paths.

3 Contract Graph Properties

In a simple version of the contract graph, each edge corresponds to a segment and is assigned a price ([Che]) but as we will see, it is not obvious what this price represents, or, in other words, which market data is used to derive the weights of edges. A more complete description of the state of a bandwidth commodity market using a contract graph, has to take into account that there is not only one price quote for each traded segment, but instead a whole set of offers is available.

Therefore we should redefine here the contract graph, in the most general case, as a graph of price quotes, not of segments only. The set of traded segments specifies the basic structure

of the contract graph, i.e. which nodes are neighbors and the number of offers for a segment determines the number of links between two neighboring nodes. Let σ be the average number of segments traded at an exchange location, and ω the average number of available contracts for a segment. Then a contract graph of ν vertices has a vertex degree equal to $\sigma\omega$ and $\frac{\nu\sigma\omega}{2}$ edges on the average, If all posted price quotes are necessary for computing path contracts, this is the graph that is needed.¹

Later in this paper we will examine some simple alternatives for compressing all the data on a traded segment to just one or a small number of useful metrics.

4 Market Data

Although trading path contracts is similar to QoS routing, and indeed a means of configuring communication paths at the inter-domain level, there are substantial differences which can be intuitively attributed to the fact that the dynamics of bandwidth trading are dictated mostly by competitive forces in the market and it is arguably not clear if and how they are linked to network dynamics, which are prevalent in the routing of data traffic. Instead of relying on network state and user traffic information to compute an adequate path between two points in a network graph, we rely on customer requirements and market data when computing a path in the contract graph. It is instructive to examine the market data which is – or could be made – available to a broker and evaluate how valuable and how costly the use of this data is.

Each broker is assumed to be connected to one or more Market Data Providers (MDP) who have access to real-time market data on several exchanges. MDPs will typically provide everyone subscribed to their information services the following information:

- *Full quotes*, i.e. complete descriptions of all bids and offers posted on all exchanges the MDP has access to.² Full quotes will typically be updated with a high frequency in the MDP's databases, but will be relatively expensive for subscribers to download due to their size and frequency of updates.
- *Spot rates*, i.e. current best offers for all segments traded on all exchanges. Update frequency is approximately the same as above, but downloading and processing all spot rates is much cheaper than full quote data. Today on-line bandwidth traders offer spot rates to the general public for free.
- *Market prices*, i.e. prices of last sales for all segments traded on all exchanges. These may or may not be disclosed, since they are part of the possibly confidential information of a bilateral agreement between the buyer and seller. If available at all, they are updated with the same frequency as spot rates and cost the same. In today's bandwidth markets

¹Judging from today's bandwidth trading and anecdotal estimates, we assume that bandwidth between only a few city-hubs will be traded in the near future, i.e. the number of vertices in the contract graph will be low for some time. On the other hand, cities with high supply and demand for bandwidth are chosen first when introducing commodity markets, so the average vertex degree of the graph will be high.

²In a perfectly commoditized and liquid market the only data of interest to traders would be the current market price for a segment contract, defined as the price of the most recent trade. Today's market, however, is not trading fully standardized contracts, nor is trading very frequent.

such data is usually not available and arguably not of much value anyway, since the number of trades is low and the speed of negotiations is much lower than the speed of the market.

- *Price indices* for selected segments. The selection of segments for indexing depends on multiple criteria, one of them being segment dependencies, explained later in this paper. Update frequency may approach that of indices in older, established commodity markets in the future.

5 Using Market Data for Brokerage

The main problem the broker is facing is that although he may have a local copy of the contract graph structure, all this graph shows is possible paths through the exchanges, i.e. which segments are being traded, not the prices for contracts on those segments, i.e. the weights on the graph are generally missing. Once a weighted graph is in place, the extended Dijkstra algorithm described in [Che] can be applied to find quality-constrained least-cost paths in that graph. In order to show that it is no simple issue to choose the right input data for the algorithm, we examine the following alternatives.

5.1 Least-Cost Graph

A first, simple option is to construct a least-cost graph (LCG) using only spot rates. Spot rates change at market speed, so the LCG would have to be updated at the same speed. Then it is a simple matter to find the shortest path in the least-cost graph. Although this *min-min* approach seems reasonable, it has disadvantages. First of all, it is overly optimistic. The shortest-path algorithm assumes the broker will be able to buy at the currently best prices for all segments and estimates the total path cost based on that assumption. In the just-in-time allocation scenario of Section 2, it is likely that when the customer decides to buy the path contract, thus leading to a series of bids by the broker in all the involved exchanges, the resulting price for the path will be higher than the estimate. The broker is committed to the path contract price and any underestimation of the path's price will directly translate to an unaccounted cost for the broker. A more subtle caveat is the possibility that the actual computed path is not actually the cheapest at the time of purchase, since no effort was made to anticipate future prices.

5.2 Highest-Cost Graph

To avoid some of the shortcomings of the first approach, we can consider a *min-max* approach. Instead of getting the spot rates (i.e. the best offers), use the worst (most expensive) offers for all segments to construct a highest-cost graph (HCG). Then find the shortest path in this graph. The benefit of this approach is that the broker can then present to the customer a "safer" cost estimate, since the price at time of purchase will be close to the estimate in the worst case. The path will probably be cheaper than estimated, so the broker will be well within the committed path contract price. But, what is the probability that the path will be even more expensive than the estimate? Lower than in the first alternative, but still

present and not taken into account. The HCG approach improves on the broker's reliability but presents us with a different problem; it is highly susceptible to manipulation. A provider wishing to keep buyers away from segments where he has not a good market position, or where extreme competition is shrinking his profit margins, publishes overly expensive offers at these exchanges. This will have the effect that the HCG-based shortest path algorithm will always avoid those segments.

5.3 Market-Price Graph

The market price at which a sale was last made is considered to be a good indication of an asset's value. A graph with market prices as weights should therefore provide a good measure of path costs. However, earlier in our discussion of market data, it was explained that this may not be the case in a low-liquidity market and why this data may not be available at all. Update frequency need not be as high as in the previous cases of this section, since trades will generally not occur as often as bid and offer postings.

5.4 Derived-Metric Graph

Since a broker may encounter problems accessing data considered sensitive, and there are pitfalls in the use of some of the above alternatives, it might be safer to depend on *derived metrics*, computed from the available market quotes. Such a simple metric could be the arithmetic average or median of all offered prices for a contract on a particular segment. Derived metrics can be computed on a daily basis, or at regular intervals during the day. The question, of course, is which type of derived data is more helpful. The spot rate Ask_{min} or arithmetic average of offers \overline{Ask} , combined with the relative offer spread ($\frac{Ask_{max}-Ask_{min}}{Ask}$) or relative bid-ask spread ($\frac{Ask_{min}-Bid_{max}}{Mid}$), where $Mid = \frac{Ask_{min}+Bid_{max}}{2}$, should provide a sufficiently informative compact view of a segment's price and risk in the short-term. Estimates of future prices and their respective deviations should be used for mid- and long-term decisions.

6 The Situation Today

Although still at a nascent stage, online spot and forward markets for bandwidth already in operation today can give an indication of requirements for and the potential of bandwidth trading in the future. A function of most web sites where bandwidth is currently being traded, is that of a bulletin board for the advertising of bids and offers. However, some are already operating real-time bandwidth exchanges (including delivery) with participation limited to a small number of carriers. Two unsolved, interrelated issues, discussed in [Enr99, dPG98], are contract standardization and market liquidity. Without the standardization (also termed "commoditization") of some bandwidth contracts, comparing offers and bidding rationally requires extensive technical knowledge of the underlying good and negotiations are slow. Directly related to this issue is the level of market liquidity, i.e. the frequency of trades.

However illiquid the current market may be, we will proceed to a preliminary study of the limited amount of real-world market data that is publicly available today. The purpose is not to produce a coherent study of the market, but rather to acquire a first impression of the

current situation, provide examples for the previous discussion on the use of market data and identify some real-world issues relevant to brokerage research.

6.1 A Market Snapshot

The data studied in this paper were collected on the 18th of December, 1999, from two well-known bandwidth trading sites located in San Francisco and London [Rat, BX] and consist of all available (posted) offers on the OTC (over-the-counter) market on that day. We assume that most, if not all, of these offers were posted in 1999, some dating a few months before the query date. Unfortunately many of them were not dated, so it is not possible to track the evolution of prices. Instead, we compiled the data into lists of destinations reachable from some “popular” cities and assumed all posted offers to be valid.

To attain comparable prices, we focused our attention to a large subset of the data, called the *focus set* in this paper. The set consists of US, European, US to Europe and US to World – destinations outside Europe – segments.³ Offers included in the set reference popular data transfer standards, namely OC-3/STM-1, DS-3, E1 and T1, which translate to the following bandwidth “grades” respectively: 155 Mbps, 45Mbps, 2Mbps and 1.5Mbps. We use the term “grades” instead of “capacities”, because the above mentioned standards refer to specific transmission and encoding or framing schemes.⁴

Most of the posted offers specify a minimum contractual term of 1 year. The long duration of these contracts is a side-effect of an illiquid, non-commoditized market. As liquidity increases and bandwidth on-demand becomes reality, more short-term contracts should appear. The monthly rates of offers which specify a longer term may include loyalty discounts, therefore these offers were excluded from the focus set. Another cause for distortion of prices is the inclusion of local loop charges, which is not always clearly indicated. This, however, is usually not the case, so any such distortion does not affect our data set significantly. QoS guarantees are usually not stated explicitly in these offers and therefore do not have a distinct, quantifiable impact on demand. We therefore assume that all offers of a particular bandwidth grade have similar quality characteristics and concentrate on prices for destination pairs.

6.2 General Observations

On the date of the data collection there were 473 offers posted on the two websites, to destinations around the world, predominantly in Europe and the United States. Out of these 397 (83.9%) met the criteria to be included in the focus set.

Convergence to a few Standards. Figure 1 provides more detail by showing the percentages of offers for each grade in the focus set. We see that DS-3 and OC-3/STM-1 offers are dominating the broadband market and we will therefore limit some of our study to these offers

³“Europe to World” segments were not included because of the very small number of posted offers of this category regarding grades in our focus set.

⁴DS-3 is wireline connectivity with a maximum sustainable bit rate of 44.7 Mbps, OC-3 (SONET) and STM-1 (SDH) are framing schemes typically for transmission over optical fiber, with a bit rate of 155.52 Mbps, T1 is a US standard for 1.5 Mbps private line circuits, whereas E1 is the European standard for 2Mbps [Min91].

later in the paper. Seeing that the market is converging to a just couple of data transport standards is an encouraging sign for bandwidth commoditization.




Figure 1: Relative volume of DS-3, OC-3/STM-1, E1 and T1 offers of the focus set, posted online. “Various” stands for all other posted offers not belonging to the focus set.

US and European Activity. Figure 2 shows the relative volume of offers in the focus set for each geographical region. It is interesting to observe that the number of bandwidth offers in Europe is as high as in the US. Relaxing our interpretation of the data, we could say that market activity levels are comparable in the two regions. Since the European and US markets dominate the world of bandwidth trading, it is to be expected that a significant part of inter-continental segments are those crossing the Atlantic.




Figure 2: Relative volume of offers of the *focus set* for each of the geographical regions of city-pair locations in the set.

Local and Global Trading. Figure 3 shows the actual numbers of offers in the focus set, grouped per grade *and* location. The most notable difference between US-internal and European quotes is the clear dominance of DS-3 bandwidth in the US, whereas OC-3/STM-1 and E1 are most common in Europe. In other words, although the numbers of offers in the two regions are comparable, there are more high capacity offers in the US. It is also interesting to note that DS-3 is not as popular for transatlantic segments, where the norm seems to be OC-3/STM-1 (more of a wholesale market).

Although we’ve identified a standards convergence trend, regional standards may continue to exist in the future, in which case we could face the definition of different bandwidth commodity contracts for the two sides of the Atlantic. Should this be the case, then the scope of bandwidth brokerage could be limited to distinct regions, unless path trading methodologies and algorithms can cope with multiple standards (assuming a global bandwidth delivery infrastructure which will not pose technical compatibility constraints).




Figure 3: Number of offers per bandwidth grade and city-pair location for the focus set.

6.3 On Prices, Graphs and Risk

In the following lines we will attempt to identify opportunities, pitfalls and research issues regarding bandwidth brokerage, by examining price quotes for particular segments.

6.3.1 Large Ask Spreads

Figure 4 shows the prices of all DS-3 offers regarding segments between New York or Washington DC (grouped as “NY/WDC”) and US destinations, where only segments with at least two posted offers are considered.⁵ Figure 5 offers the same picture for Los Angeles and San Diego (grouped as “LA/SD”).⁶ When considering the total set of offers for a particular segment, we

⁵Lower prices generally correspond to more recent offers, due to the strong downward pressure on prices in the bandwidth market.

⁶State names indicate the grouping of offers to several destinations in a state.

observe in some cases large deviations from the average, in the order of a few thousand dollars per month. This simply means that in cases where a trader wishes to buy more than one contract on a segment, he should clearly not take only the price of the best offer into account. This also holds for trading a single contract, because in the case where a buyer misses the best offer, acquiring the second best may incur a substantial additional cost. These spreads are a sign of illiquidity *and* incomplete commoditization. In a liquid commodity market trading non-differentiable standard contracts competition would force asks to concentrate much closer to the market price.

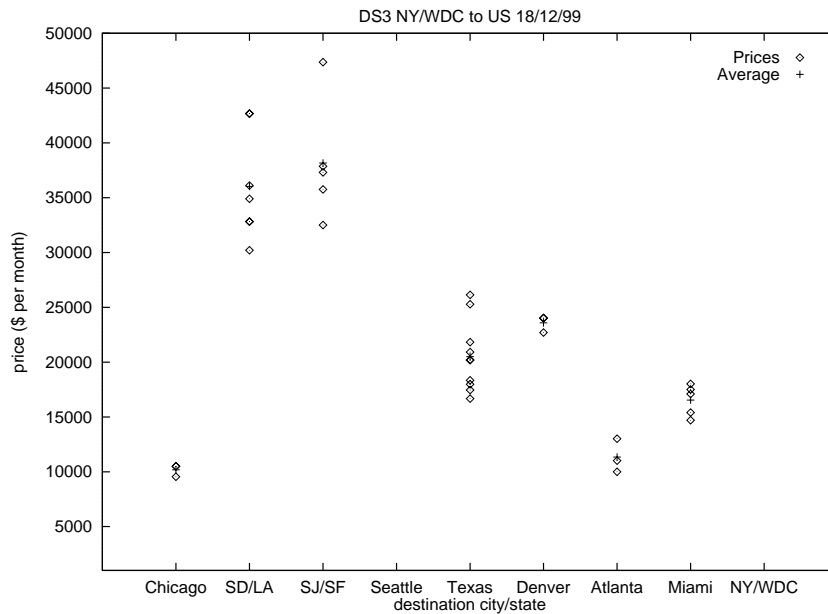


Figure 4: DS-3 prices from New York and Washington DC to the rest of the US.

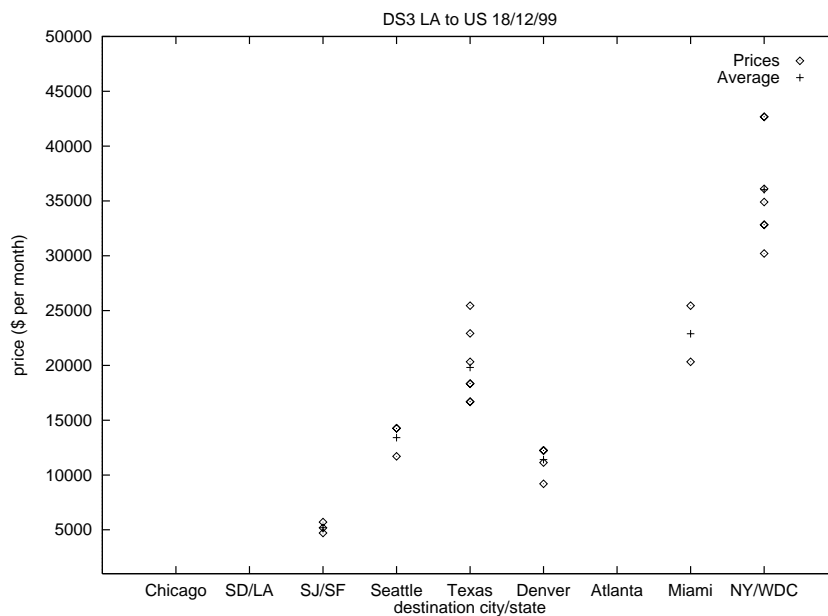


Figure 5: DS-3 prices from Los Angeles and San Diego to US destinations.

It is clear then that besides the long-term risk of investment in a commodity, due to changes in supply and demand, there is considerable short-term risk in bandwidth trading, due to low liquidity and the lack of fully commoditized contracts. Risk becomes even more important in path trading, because the selection of a path in the contract graph depends on the prices of multiple segment contracts, each having a degree of uncertainty associated with it. To quantify short-term risk in the most simple manner, we measure the ask spread of each segment, as defined in section 5.

By combining Figures 4 and 5, one can attempt to identify arbitrage opportunities for the NY/WDC-LA/SD pair. However, pure arbitrage (i.e. no risk) is rare and indeed including risk in the decision process is very valuable, as the following example clearly illustrates. When examining only the best offers for each segment, it appears to be slightly more expensive to buy a path through Denver (\$31,900) in place of a direct New York to LA offer (\$30,200). Surprisingly though, the average price for a path through Denver is \$34989, \$1036 lower than the average direct NY-LA price. On top of that, NY-Denver-LA path price quotes have a much lower spread. Consequently, in cases where more than one contract are needed or in the just-in-time allocation scenario, or when it is not certain that the best offer is valid and will be delivered, the path NY-Denver-LA may be preferable and in any case should be taken into account.

Returning to the discussion of contract graph options, in Section 5, the existence of short-term risk reflected in large ask spreads renders the highest-cost graph useless (unreliable) and favors the derived-metric and least-cost graphs (or a combination thereof). In the future, when bandwidth commodity markets mature, market-price graphs will be more useful, enhanced with mid- and long-term forecasts and risk measures.

6.3.2 Segment Dependencies

Figure 6 combines STM-1/OC-3 offers from London and New York to European destinations in the same plot. Country names are used when prices for several destinations in a country are grouped together.⁷ By observing the average prices from London and New York to all destinations on Figure 6, it becomes clear that there is a dependency. In particular, those destinations which are relatively expensive (cheap) to reach from London, are also expensive (cheap) to reach from New York. This is no surprise, since many transatlantic connections are set up through London and therefore NY-London is a hidden component of many NY-Europe offers. The topology of underlying networks *is* therefore important for understanding the commodity market, inasmuch as it causes stochastic dependencies among certain segments.

Such dependencies point out the value of the creation of bandwidth price indices for those segments which underlie many offers (in this case NY-London), since changes in the prices of these segments will have a large overall impact on the market.

The existence of dependencies does not in any way preclude the possibility of geographical arbitrage on stochastically dependent segments. As an example, if we take just the average prices of Figure 6 and compare them we see that a buyer interested in NY-Europe segments can get as much as 4% (\$8,900 for NY-DK/SE), in the average, off the monthly rate when including NY-London-Europe paths in the selection.

⁷The two-letter country codes used in Figure 6 correspond to internet national domain abbreviations (ISO 3166).

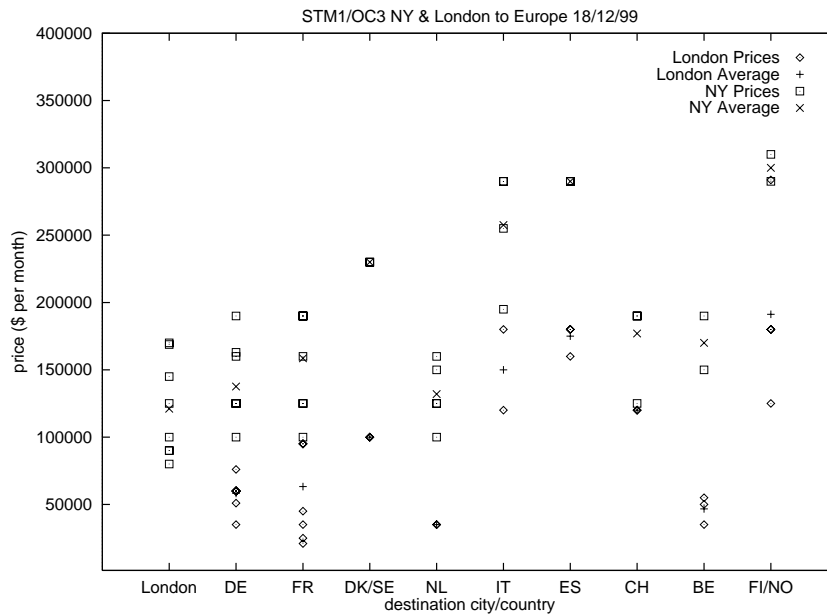


Figure 6: STM-1/OC-3 prices from New York and London to European destinations.

7 Conclusions

It was shown that in today's illiquid and only partially commoditized bandwidth market a bandwidth broker is facing considerable risk, especially when trading composite "path" contracts. This was supported with real-world examples. Conditions for the management of risk are the formulation of uncertainty concerning prices and contract availability and the inclusion of risk measures in the decision-making process and supporting tools of the broker. In particular, a broker's trading algorithms should not rely just on spot rates, but rather make use of statistical metrics, even in the short-term. The observation and interpretation of large spreads in a segment's price quotes and stochastic dependencies among segments are helpful in modelling the bandwidth commodity market. As a side-remark, it seems traders may need to deal with the existence of regional capacity standards in the future, although a couple of global standards should eventually prevail. Clearly more research is needed in the field of bandwidth brokerage, to understand the workings of a new market and support rational trading decisions which will hopefully lead to efficient allocation of communication resources between network domains.

References

- [BX] Band-X. Homepage, <http://www.band-x.com>.
- [Che] G. Cheliotis. A market-based model of bandwidth and a new approach to end-to-end path computation with quality guarantees. Presented at the MIT/Tufts Workshop on Internet Service Quality Economics (ISQE), 2-3 Dec. 1999, Cambridge MA. <http://www.marengoresearch.com/isqe/>.

- [CL00] G. Cheliotis and B. Liver. Brokerage of bandwidth commodities. In *7th International Conference on Intelligence in Services and Networks, 23-25 Feb. 2000, Athens, Short Paper Proceedings*, 2000.
- [CMMG97] D. Cooperstein, C. Mines, M. MacDonald, and M. Goldberg. Bandwidth's new economics. Forrester Research, 1997.
- [CN98] 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. <http://www.enron.net>.
- [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.
- [IFC] IFCI. International finance and commodities institute, risk watch homepage, <http://risk.ifci.ch>.
- [Min91] D. Minoli. *Telecommunications Technology Handbook*. Artech House, Norwood, MA, 1991.
- [Rat] RateXchange. Homepage, <http://www.ratexchange.com>.
- [SB99] A. Segev and C. Beam. Brokering strategies in electronic commerce markets. In *E-COMMERCE 99, Denver, Colorado*. ACM Press, 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.