

Market-Based Call Routing in Telecommunications Networks using Adaptive Pricing and Real Bidding

M.A.Gibney¹, N.R.Jennings², N.J.Vriend³

¹Department of Electronic Engineering, ³ Department of Economics,
Queen Mary and Westfield College, University of London,
London E1 4NS, UK.

² Department of Electronics and Computer Science
University of Southampton,
Highfield, Southampton SO17 1BJ, UK.

1 3 {M.A.Gibney, N.Vriend} @qmw.ac.uk

2 nrj@ecs.soton.ac.uk

Abstract

We present a market-based approach to call routing in telecommunications networks. A system architecture is described that allows self-interested agents, representing various network resources, potentially owned by different real world enterprises, to co-ordinate their resource allocation decisions without assuming *a priori* co-operation. It is argued that such an architecture has the potential to provide a distributed, robust and efficient means of traffic management. In particular, our architecture uses an adaptive pricing and inventory setting strategy, based on real bidding, to reduce call blocking in a simulated telecommunications network.

1 Introduction

In telecommunications networks, call traffic is typically routed, through the network from source to destination, on the basis of information about the traffic on that path only. Therefore, path routing is carried out without regard to the wider impact of local choices. The main consequence of this myopic behaviour is that under heavy traffic conditions the network is utilised inefficiently: rejecting more traffic than would be necessary if the load were more evenly balanced. One means of performing such load balancing is to centrally compute optimal allocations of traffic across the network's paths using predictions of expected traffic (Bertsekas and Gallager 1987). When such calculations have been completed, the network management function can configure the network's routing plan to make the best use of the available resources given the predicted traffic. However, as networks grow larger and involve more complex elements, the amount of operational data that must be monitored and processed (by the network management function) increases dramatically. Therefore in centralised architectures, management scalability is bounded by the rate at which this data can be processed (Goldszmidt&Yemini 1998). In addition, there are a number

of well known shortcomings with algorithms to compute optimal network flows; these include progressively poorer performance in heavily loaded networks and unpredictable oscillation between solutions (Kereshbaum 1993). Furthermore, the very centralisation of the network management function provides a single point of failure; thus making the system inherently less robust.

For the above mentioned reasons, a decentralised approach to routing is highly desirable. In such cases, decisions based on more localised information are taken at multiple points in the system. The downside of this, however, is that the local decisions have non-local effects. Thus, decisions at one point in the system affect subsequent decisions elsewhere in the system. Ideally localised control would take place in the presence of complete information about the state of the entire system. Such a state of affairs would enable a localised controller to know the consequences of a choice for the rest of the network. However, there are two main reasons why this cannot be realised in practice. Firstly, the network is dynamic and there is a delay propagating information. This means that a model of the network state held at any one point is prone to error and difficult to keep up to date. Secondly, the scaling issues involved in making flow optimisation computations for the entire network (noted above) would obtain here also. Therefore, a system in which local decision making takes place in the presence of an incomplete view of the wider network is the only feasible solution for providing distributed control.

A promising approach that combines the notion of local decision making with concerns for the wider system context is that of agent technology (Jennings and Wooldridge 1998). Agents address the scaling problem by computing solutions locally, based on limited information about isolated parts of the system, and then using this information in a social way. Such locality enables agents to respond rapidly to changes in the network state; while their sociality can potentially enable the wider impact of their actions to

Copyrighted Material

be coordinated to achieve some socially desired effect. Systems designed to exploit the social interactions of groups of agents are called multi-agent systems (MAS). In such systems, each individual agent is able to carry out its tasks through interaction with a small number of acquaintances. Thus, information about the extent of the system is distributed along with whatever functionality the MAS is designed to perform.

One agent-based technique that is becoming increasingly popular as a means of tackling distributed resource allocation tasks is market-based control (Clearwater 1996). In such systems, the producers and consumers of the resources of a distributed system are modelled as the self-interested decision-makers described in standard microeconomic theory (Varian 1992). The individual agents in such an economic model decide upon their demand and supply of resources, and on the basis of this the market is supposed to generate an equilibrium distribution of resources that maximizes social welfare. In market-based control, the metaphor of a market economy as a system computing the behaviour that solves a resource allocation problem is taken literally and distributed computation is implemented as a market price system. That is to say, the agents of the system interact by offering to buy or sell commodities at given prices (Wellman 1996). In our case, such an approach has the advantage that ownership and accountability of resource utilisation are built into the design philosophy. Thus, market-based solutions can be applied to the management of multi-enterprise systems without forcing the sub-system owners to co-operate on matters affecting their own commercial interest.

Within this context, this paper describes a system to balance traffic flow through the paths of a logical network, based on the local action of agent controllers coupled with their social interaction as modelled by a computational market. The approach presented here builds upon the preliminary work reported in (Gibney and Jennings 1998) in that it shares the same architecture, roles and deployment model for the agents. However, to improve upon our earlier results we devised a new approach to the way that agents adapt their pricing and inventory strategies according to the outcome of individual market actions and the profitability of trading in the market. More generally speaking, this paper extends the state of the art in market-based control in the following ways. Firstly, it models a complex two-level economy, in which not only end users but also the internal components of the system compete with one another for resources. The rationale behind using a two-level economic model is to realise call admission control in the same framework as the network management function. This is novel, as market-based control has not previously been used to address two control issues in the same system. Particularly, having two kinds of market within the economy, with agents active in different roles in each of them, provides an elegant way to acquaint agents with one another. This architecture also provides an appropriate way to situate the intelligence of the system in a multi-enterprise

network with self-interested enterprises. Secondly, a novel approach is adopted to pricing strategy. Our agents adapt to the outcomes of market interactions which use real bids and offers (i.e. agents state a price in an auction-like market and are then committed to buying or selling at that price in that session). This approach was adopted because it eliminates the lengthy series of interactions between agents that is required to calculate the equilibrium price in the market. Rather, we use real bidding and allow the agents to adapt their bidding behaviour to the outcomes of the auctions over time. Real bidding allows us to use more rapid (one shot) auction protocols as markets.

The remainder of the paper is structured in the following manner. We discuss the background and motivation for this work in Section 2. Section 3 describes the architecture of the system as a whole and the institutional forms of the possible interactions between agents. The design of individual agents is given in Section 4. Section 5 discusses the experiments carried out to evaluate the performance of the system. Finally, Section 6 details our conclusions on the work presented in this paper and discusses the open issues and future work.

2 Background and Motivation

Decentralized approaches to routing, usually in packet switched telecommunications networks, based on the interaction of controllers distributed through the network have existed for some time (Schwartz and Stern 1989). However, agent based approaches extend this idea by modelling the interaction of distributed controllers as a social process. A number of agent based solutions have been proposed to the problem of load balancing in telecommunications networks. (Appleby and Steward 1994) make use of mobile agents roaming the network and updating routing tables to inhibit or activate routing behaviours. (Schoonderwoerd et.al.96) extend and improve this approach by using ant-like mobile agents that deposit "pheromones" on routing tables to promote efficient routes (the majority of ants use the efficient routes and the pheromones re-enforce this behaviour in other ants). (Hayzelden and Bigham 1998) employ a combination of reactive and planning agents in a heterogeneous architecture to reconfigure route topology and capacity assignments in ATM networks. All of these systems exhibit increased robustness and good scaling properties compared to centralised solutions. Indeed, in network environments with symmetric traffic requirements, ant-like agent solutions have even been shown to provide superior load balancing to both statically routed networks and a more conventional mobile agent approach (Schoonderwoerd et.al.96).

However, all the aforementioned approaches model networks as a single resource and therefore act to optimise utilisation of that resource. This makes sense because a poorly managed telecommunications network benefits no one. However, the main disadvantage of such a perspective is that different telecommunications operators join

their networks together (something which is an increasingly common trend). In such cases, if the different sub-network owners agree on a single, unified static network management policy, it is unlikely that this policy will benefit all their interests individually as well as collectively over time. We address this issue by modelling our agents as the resources and groups of resources that enterprises might own or lease in a multi enterprise environment.

Another increasingly common aspect of modern telecommunications deployment is the practice of enterprises in other sectors (banking and other traditional consumers of telecommunications services) leasing bandwidth from telecommunications providers to create virtual private networks in the short, medium and long term (Cisco 1999). Again, this promotes the creation of multi-enterprise networks. In such environments each enterprise clearly has an incentive both to see that overuse does not degrade network performance and to make the greatest possible use of their network ownership. Since these parties cannot agree each traffic policy decision individually, conflicting incentives must be reconciled outside the traffic management domain. Typically this is achieved by allowing sub-network owners to set policy within the remit of their own resources (Stallings 1997). However, the static nature of these policies and the conflict between them at sub-network interfaces often causes institutionalised under-use of the network as a whole.

Both of these trends suggest that telecommunications network management, once a centralised and monolithic undertaking, will increasingly benefit from an open, robust, scalable and inherently multi-enterprise approach. Therefore, one of the aims of this work is to use the multi-agent system paradigm to address the problem of multi-enterprise ownership of the network, while simultaneously addressing the problems of robustness and scalability. Against this background, the resource allocation problem in a network with multiple, non co-operating enterprises can be recast as the problem of reconciling competition between self-interested, information-bound agents. We conjecture that a market economy might be an effective mechanism for achieving this goal. Therefore we decided to implement our telecommunications network management framework using economic concepts and techniques.

3 System Architecture

The overall system architecture consists of three layers. The lower layer is the underlying telecommunications infrastructure. The middle layer is the multi-agent system that carries out the network management function. The top layer is the system's interface to the call request software. More details of the rationale to this design are given in (Gibney and Jennings 1998). The remainder of this section concentrates on the agent layer: describing the main components (section 3.1) and how they relate to one another (section 3.2).

3.1 The Agents and their Interactions

The system makes use of three agent types: (1) the *link agents* (section 4.1) that represent the economic interests of the underlying resources of the network, (2) the *path agents* (section 4.2) that represent the economic interests of paths across the network and (3) the *call agents* (section 4.3) that represent the interests of callers using the network to communicate. A *link agent* is used for every link in the network and is deployed at the entry node for that link. A *path agent* is used for each logical path in use across the network and is deployed at the source node for that path. Here we use three path agents for each source destination pair. There is a reasonable number of alternate paths across which to share a single traffic requirement (alternate static routing systems commonly using three or fewer paths). A *call agent* is used for each source destination pair in the network and is deployed at the source of the traffic requirement that it represents.

The agents communicate by means of a simple set of signals that encapsulate offers, bids, commitments, and payments for resources. We couple the resources and payments with the offers and bids respectively. This reduces the number of steps involved in a transaction (committing agents to their bids and offers ahead of the market outcome), and so increases the speed of the system's decision making (an important consideration in this domain). To enforce these rules the interactions between the different agent types are mediated by means of market institutions (described in section 3.2).

An important notion in agent technology is that agents should be proactive (i.e. be able to anticipate the requirements of the environment and behave accordingly). In our system, we apply this concept to implementing a call routing mechanism that does not need to examine the network state before routing each call. Our path agents proactively determine how many calls they will be able to handle in advance, and seek to obtain the necessary resources to handle them. To be able to offer resources to callers pro-actively, the path agents lease bandwidth from the link agents over a period of time, paying installments on the lease at prices agreed on the link markets.

3.2 Market Institutions

Our system makes use of two types of market institution: At the *link market* (section 3.2.1), slices of bandwidth on single links (the fundamental resources of the system) are sold to path agents. At the *path market* (section 3.2.2), the slices of bandwidth on entire paths across the network are sold to call agents to connect calls.

3.2.1 Link Markets

Link markets are sealed bid double auctions. In this protocol the auctioneer receives two sets of prices in each trading period: bids for resources from buyers and offer prices from sellers, and computes the successful trades according to a set of rules defined for the auction. A sealed bid protocol was chosen because it provides a means to

Copyrighted Material

complete institutionally mediated bargaining in one shot. Therefore bargaining that would take an indeterminate time using iterated market institutions such as continuous double auctions can be completed almost instantaneously.

The resources exchanged at the *link markets* are the right to use slices of bandwidth on individual links, which when taken together, provide the necessary bandwidth to connect calls across paths. The link markets use a sealed bid double auction in which buyers and sellers periodically submit bids for individual units of the resource. In our case, buyers and sellers are constrained to their roles in the market by their position in the network. Thus, *path agents* need to buy resources from *link agents* to offer services to callers. The bids and offers are ordered from high to low, and low to high respectively. There will be a range of prices for which the market will clear at which the maximum amount of resources can be traded. Buyers bidding above these prices and sellers offering below it are allowed to trade. The buyers and sellers within this group are matched randomly so that the benefits of trade are assigned between successful agents randomly also. The trading price for each given transaction is determined at random in the range between the buyer's bid and the seller's offer. Notice that this procedure implies that no buyer will pay more than his bid, and no seller will get paid more than his offer. Moreover, the procedure implies that the total surplus realized in the market, a measure of the social welfare, is maximized because the benefits of trade are distributed randomly between all successful agents.

3.2.2 Path Markets

The path market is also a sealed bid auction. This is because it is a critical performance requirement of the system that the allocation of call traffic to paths occur almost instantaneously (so that callers are not kept waiting for calls to be established). This means the auction protocol has to be as short as possible. As before, the most efficient protocols in this respect are the single shot, sealed bid types. Since we have a single caller and multiple path agents offering resources, a single sealed bid auction is appropriate.

A buyer sending a service request message to the market initiates the auction. The auctioneer then broadcasts a request for offers to all agents able to provide the connection. All sellers simultaneously submit offers and the lowest one wins the contract to provide the connection. In this market, we experimented with two protocols: (i) The First Price protocol, in which the price at which the buyer and the seller trade is that of the highest bid submitted. (ii) The Vickrey, or second price auction protocol, in which the price at which the buyer and seller trade is that of the second highest bid submitted. We choose to experiment with two strategies because economic theory predicts that Vickrey auctions provide more competitive market outcomes, doing away with wasteful speculation by encouraging truth telling behaviour on the part of the participants (Varian 1995). However, since we are using simple adaptive agents without speculation,

we were unsure as to whether this factor would impact the overall behaviour of the system. To test the impact of this factor on the system as a whole, we implemented the market with both protocols and empirically tested the efficiency of each (section 5).

4 Designing Economically Rational Agents

The range of potential interactions is determined by the market institutions in which the agents participate. In both of our market types, agent communication is restricted to setting a price on a single unit of a known commodity. Therefore, agents set their prices solely on the basis of their implicit perception of supply and demand of that commodity at a given time. When a resource is scarce, buyers have to increase the prices they are willing to bid, just as sellers increase the price at which they are willing to offer the resource (*mutatis mutandis* when resources are plentiful). Here, agents perceive supply and demand in the market through the success or otherwise of bidding at particular prices.

4.1 Link Agents

A *link agent* has a set of n resources, slices of bandwidth capacity required to connect individual calls, that it can sell on the *link market*. At time t , the price to be asked for each of these units is stored in a vector $p_t = \{p_t^1, \dots, p_t^n\}$ with the range of possible prices being zero to infinity, $p_t^i \in [0, \infty)$ for each member of the vector $i = 1, \dots, n$ and each time period t . At time $t = 0$ the prices for each unit are randomly (uniformly) distributed on $[0, H]$ where H is the initial upper limit on prices asked. When x units have been allocated, the remaining $n - x$ units are offered to the link market for sale simultaneously. Suppose that of the $n - x$ units offered for sale in a given period t , the m units with the lowest prices are successfully sold. The prices in the vector are updated as follows:

$$p_{t+1}^l = p_t^l, \text{ for } l = 1, \dots, x \quad (1)$$

$$p_{t+1}^i = p_t^i \times (1 + \epsilon), \text{ for } i = x + 1, \dots, x + m + 1 \quad (2)$$

$$p_{t+1}^i = p_t^i \times (1 - \epsilon), \text{ for } i = x + m + 2, \dots, n \quad (3)$$

$$\text{Where } \epsilon = U(0, \sigma) \quad (4)$$

Thus the link agent increases or decreases the price of any unit by a small amount ϵ after each auction. Here ϵ is obtained from a uniform random distribution between zero and σ (here 0.1). If previously allocated units are released by the path agent, they join the pool of unallocated units and the price vector is re-ordered to reflect this. This approach was chosen so that the prices of each resource on the link, taken together, should adapt to the demand on the network to carry traffic.

4.2 Path Agents

A *path agent* acts as both a buyer of link resources and a seller of path resources. Their buying behaviour is detailed in section 4.2.1, and their selling behaviour in section 4.2.2. In general, path agents wish to buy resources cheaply from link agents and sell them at a profit to end consumers. To do this, they bid competitively to acquire resources that they then sell on to callers, at a price not less than that paid for them. The path agent tries to maximize its profits by adjusting its inventory and sales behaviour on the basis of the feedback it receives from the market. The mechanism by which the path agent decides what resource level to maintain is described in section 4.2.3.

4.2.1 Buying Behaviour

A *path agent* actively tries to acquire resources (units of link bandwidth needed to connect a call), across the chain of links that it represents. It does this by placing bids at each of the link markets at which the resources it needs are sold. The agent retains a vector of prices that it is willing to pay for resources on each of the links that constitute its path. The agent's strategy is to try to equalise its holding of resources across each of those links; uneven resource holdings have to be paid for but cannot be sold-on or bring in any revenue because they do not constitute complete paths. The *path agent* tries to maintain its resources at a level w that is discovered through hill climbing adaptation (Russell&Norvig 1995) to the behaviour of the market (section 4.2.3). The most profitable value of w is obtained by adjusting it according to changing profit during ongoing buying and selling episodes. When x units have been acquired, the path agent bids for the remaining $w - x$ units on the link market simultaneously. Suppose that, for a given link, of the $w - x$ units that the path agent bids for at time t , the m units with the highest prices are successfully acquired. The prices in the vector are updated as follows (using ϵ as defined previously in section 4.1).

$$p_{t+1}^i = p_t^i, \text{ for } i = 1, \dots, x \quad (5)$$

$$p_{t+1}^i = p_t^i \times (1 - \epsilon), \text{ for } i = x + 1, \dots, x + m + 1 \quad (6)$$

$$p_{t+1}^i = p_t^i \times (1 + \epsilon), \text{ for } i = x + m + 2, \dots, w \quad (7)$$

This price setting mechanism was chosen because it allows the path agents to adaptively determine prices for individual link resources. The price bid for each resource should be as low as possible without failing to win the resource in the auction. Therefore the agent makes a bid for each resource that it needs separately. If a bid fails, the agent increases the price it will bid at the next auction (in order to increase its likelihood of winning the resource). If a bid succeeds, the agent reduces the price it bids for that resource in subsequent auctions (in order to avoid paying over the market price).

4.2.2 Selling Behaviour

A *path agent* will offer to sell a path resource whenever an auction is announced by the path market and it has an appropriate path resource to sell. The price asked is determined by the cost of acquiring the underlying link resources and the outcomes of previous attempts to sell. Let p_t be the price of a path resource at time t (the time of the auction) ranging from the cost to acquire the resource to infinity, $p_t = [\text{cost}, \infty >$. The first time the agent offers a resource for sale, it offers it at a price given by $p_t = \text{cost} \times (1 + \epsilon)$ in order to sell at a profit (using ϵ as defined previously in section 4.1). Subsequently the offer price is given by:

$$p_{t+1} = \text{Max}(p_t \times (1 + \epsilon), \text{cost}) \text{ if last offer was successful} \quad (8)$$

$$p_{t+1} = \text{Max}(p_t \times (1 - \epsilon), \text{cost}) \text{ if last offer was not successful} \quad (9)$$

This price function ensures that the path agent never sells a resource for less than it paid to acquire it in the first place. Given its inventory level (see section 4.2.3), the agent attempts to maximise its income. The price bid for each resource should be as low as possible without failing to win the resource in the auction. Therefore, the agent increases the price it asks whenever it is successful, and decreases it whenever it is unsuccessful. This means the agent adapts its price to the level of competition, as it perceives it from the outcomes of previous auctions.

4.2.3 Inventory Level

The resource levels of the various path agents determine the maximum flows available for traffic on individual paths. In our case, the system design philosophy is to have the individual paths determine their own optimal resource levels. When this is achieved, balancing the load in the network as a whole becomes an emergent property of the social interactions of the agents. Path agents act in response to the economic pressures exerted on them by their consumers, competitors and suppliers. Therefore, we choose to have path agents discover their own optimal flows by adaptation to economic conditions as they perceive them (through interactions in the markets in which they compete). Path agents are both *buyers* and *sellers* that attempt to maximise their profit through trade, where profit is the difference between their revenue (from selling path resources to callers) and their expenditure (cost of acquiring and holding onto resources). In order to maximise profit agents must have an inventory level that is optimal for them, in the competitive environment in which they find themselves. Therefore our path agents adapt their inventory levels to the profits they earn through their interactions with the market. This is implemented by having the inventory level of individual path agents climb the hill

of their profits.

In more detail, let R_t be the profit of an agent at time t and I_t be the desired resource inventory of that agent. If profit has increased since the last market interaction ($R_t > R_{t-1}$) and the last change corresponded to an increase in desired inventory level ($\Delta I_t > 0$), the new desired inventory level is increased by one resource unit. If the last change in desired inventory was negative ($\Delta I_t < 0$) then desired inventory is reduced by one unit. However, if profits have fallen, ($R_t < R_{t-1}$) and the last change was positive ($\Delta I_t > 0$) we decrease the desired inventory level; If the last change was negative ($\Delta I_t < 0$) then the desired inventory level will be increased. When decreasing the desired inventory level, the agent chooses to give up the most expensive of its link resources (that are not allocated to a call). This strategy reduces the agents inventory rental cost by the largest amount possible in a single time step.

4.3 Call Agents

Call agents initiate the auctions at which path agents compete by signalling that they wish to buy resources to a given source destination pair. In doing so they set a maximum price that they will not exceed to make a call. This puts downward pressure on the offers made by path agents to provide resources across whole paths, thereby anchoring the system.

5 Experimental Evaluation

Our experiments were designed to test three hypotheses. Firstly, we wanted to know whether market-based systems can compete with static routing algorithms in terms of call routing performance (section 5.1). Secondly, we wanted to know if our system uses the network efficiently (i.e. does it use the best routes) whenever possible, allowing for congestion (section 5.2). Thirdly, we wanted to test if the system discriminates in its choice of routes between paths that would be indistinguishable from one another to a conventional static routing algorithm, without expected traffic predictions (i.e. they differ only in their proneness to congestion) (section 5.3).

5.1 Performance Evaluation

In terms of performance, we sought to address two fundamental questions. Firstly, can our market-based control system perform as well or better than a conventional system? Secondly, what is the effect of using a Vickrey auction protocol rather than a first price auction protocol at the path market?

In a series of experiments we tested the efficiency of our market-based control mechanism (using both first price and Vickrey auctions as path markets) against a static routing mechanism. The static routing mechanism used a number of paths (three) between each source destination pair and routed traffic to these in order of path length, using longer paths when the shorter ones became congested. Here

efficiency was measured as the proportion of calls successfully routed through the network as a percentage of the total number of calls offered. The experiment was configured to simulate a small irregularly meshed network of 8 nodes with link capacities sufficient for 200 channels. Calls arrived on average approximately every 5 seconds, were routed between a randomly chosen source destination pair, and lasted an average of 200 seconds. Call arrival and call duration were determined by a negative exponential time distribution function, with $U(0, 1)$ being a uniform random distribution between 0 and 1. The inter-arrival time between calls and the call duration were calculated using the formula: $f(x) = -\beta \ln U$, where β was the mean inter-arrival time and mean call duration respectively. The simulation was allowed to run for 20,000 seconds in each case. The traffic model was chosen to simulate a realistic call arrival rate and duration. The network dimensions were chosen to reflect a small network under heavy load.

These results show that similar levels of performance are obtained using the market-based control mechanism (95.4% of calls connected) and static routing (94.8% of calls connected) (Figure 1.). It is interesting to note that, contrary to our original hypothesis, using a Vickrey auction for the path market did not improve upon the results obtained using a first-price auction. One possible reason for this is the way in which the path agents in our system adapt their pricing strategy to market outcomes.

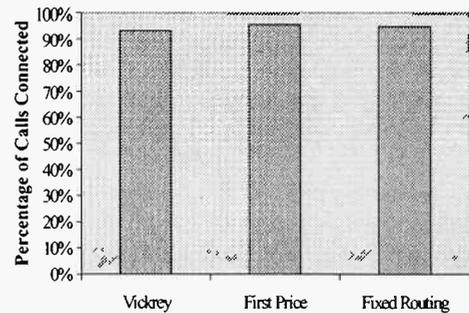


Fig. 1. Performance of Market-Based and Static Routing

Vickrey auctions are designed to make the markets more efficient by making counter speculation between competing agents wasteful. However, with simple adaptive agents being used here such speculation does not occur, so the effect should not be significant.

The ability of our system to perform as well as a static routing mechanism should be taken as a positive result. As well as matching the performance of conventional routing techniques, our market-based approach has a number of distinct advantages for network operators and users. Firstly, it provides an architecture that is open to deployment in multi-enterprise environments without the inefficiency of static internetworking policies at sub-network interfaces. Secondly, our system is scalable in that no agent has to know the address of a significant number of peer agents or possess a map of the entire network. Thirdly, our system

allows a much quicker response to call requests because the call routing process does not need to obtain information from the wider network at call set up time. In our case a call request can be processed and the call dispatched to a path (or refused) solely on the basis of information present at its source. This is achieved by having path agents that pro-actively determine their capacity to handle traffic, rather than waiting for call requests before processing.

5.2 Resource Utilisation Efficiency

In addition to the raw connection performance, it is important to know how effectively the network's resources are used. This is important because both the callers and the system benefit from routing calls through the shortest path when one is available. Shorter call routes use fewer system resources than longer ones, and they provide a service with less delay to end users. To assess our system's performance, we analysed the relative percentages of calls that were assigned to first, second and third best paths (by the number of links which make up the path). The data (Figure 2.) shows that the majority of calls were routed through the most efficient routes: 61% using the most efficient route, 25% the second best route and 15% the third best. Thus not only does our system route most of the calls it is presented with, it also makes the most efficient routing choices.

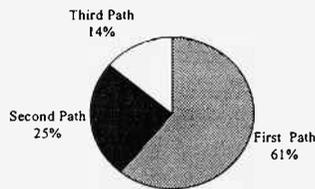


Fig. 2. Utilisation of First, Second and Third Shortest Paths by Market-Based Control

5.3 Congestion Discrimination

One of the claims we made for market-based approaches, is that good system level choices can emerge from local choices, that are influenced by information about the social context (obtained through interaction). To explore this hypothesis, we examined the performance of our system in cases that are indistinguishable from a local perspective. Thus, we focussed on source-destination pairs where all the routes are of equal length. In such cases, an alternate routing mechanism cannot decide between these paths without some notion of congestion through the whole network, which cannot be calculated and propagated in real time (recall the discussion of section 2). Alternate routing mechanisms can assign traffic to paths probabilistically, so that statistically, over time more traffic is routed to less congestion prone paths. However, this method is dependent on the accuracy of past measured traffic as a predictor of future traffic patterns. With our approach, we believe such

discrimination would emerge from the competitive nature of the market place. The reason for this hypothesis is that while path agents for paths of equal route length have to obtain the same number of resource slices, some have to obtain the more congestion prone of those slices. By definition, the more congestion prone resource slices are traded in the more competitive (and hence more highly priced markets). All other things being equal, the profitability of selling these paths will be lower because of the higher costs. With lower profitability comes a lower inventory level and fewer calls being routed via that path.

We wanted to test whether the market-based control mechanism is able to discriminate between paths on the basis of congestion cost in real time. In order to investigate this, we examined the source-destination pairs in our network configuration (nine of them) for which all (three) known paths consisted in an equal number of links. We plotted the percentage of calls routed to each of the three paths in (reverse) order of congestion and took the average of these values across all nine source destination pairs (Figure 3.). Our results clearly show that the market mechanism is able to distinguish between congestion costs entailed in routing across paths of otherwise equal length and assign calls to the least congested path most of the time.

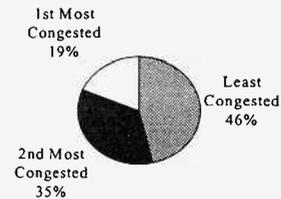


Fig. 3. Utilisation of paths in order of congestion by market-based control

6 Conclusions

We have described the design and implementation of a market-based system for call routing in telecommunications networks. Our system performs comparably with a static routing approach in terms of the percentage of the calls that are connected. However, from an architectural point of view, the market-based approach represents an improvement on static, centralised systems for a number of reasons. Firstly, it provides a platform for implementing network traffic management in a multi-enterprise internetwork. Secondly, it does not rely on a centralised controller to compute network reconfigurations, making the network management function robust to failure. This means that the system described in this work is architecturally more robust than an equivalent network with a centralised control mechanism. It is important to distinguish between this sense of robustness, and robustness as an empirical measure of the performance of the system in the event of

Copyrighted Material

the failure of a component of the system (or the agent process managing it). The performance of a decentralised control system degrades with the loss of controlling processes, while a centrally controlled one would be without control if the controlling process were to be lost. The study of the degradation of control efficiency in the system described here has been left to future work. Thirdly, no agent needs to know of the existence of more agents than there are links in the paths of the network (making the agent acquaintance databases compact and the whole system more scalable). Fourthly, there is no requirement to test the network state at call set up time, making the call set up procedure faster and more robust. Fifthly, the cost of each call to the network and the proportion of that revenue owing to each of the enterprises involved in carrying that call can easily be computed from information available to the user terminal equipment when the call is made, thus making call charging more efficient.

The results presented in this paper show that our market-based system performs the call routing and network management tasks adequately. However, the function used to determine the inventory level of path agents is quite simple and responds reactively to burstiness in the call arrival rate; this may be inducing unwanted oscillation in the path inventory parameter which may be adversely affecting performance. We intend to experiment with this function, and the parameters that govern its behaviour, to determine the impact of our choices on the performance of the overall system, and to see if that performance can be improved upon.

Acknowledgement

This work was carried out under EPSRC grant No.GR/L04801.

References

- Appleby, S., and Steward, S. 1994. Mobile Software Agents for Control in Telecommunications Networks, *BT Journal of Technology* 12 – 2: 104 - 113
- Bertsekas D. & Gallager M. 1987. *Data Networks*, Prentice Hall International, Inc.
- Cisco Systems Inc., 1999. A Primer for Implementing a Cisco Virtual Private Network (VPN), Cisco Systems Inc.
- Clearwater, S. H. 1996. *Market-Based Control A Paradigm for Distributed Resource Allocation*, World Scientific Press.
- Gibney, M.A. & Jennings, N.R., 1998. Dynamic Resource Allocation by Market-Based Routing in Telecommunications Networks, In *Proceedings of Intelligent Agents for Telecommunications Applications 1998*, Springer Verlag. 102 - 117
- Goldszmidt, G. and Yemini, Y., 1998. Delegated Agents for Network Management. *IEEE Communications Magazine* March 1998: 66 -70
- Hayzelden, A. and Bigham, J. 1998. A Heterogeneous Multi-Agent Architecture for ATM Virtual Path Networks, In *Proceedings of Intelligent Agents for Telecommunications Applications 1998*, Springer Verlag 45 - 59
- Jennings, N.R. & Wooldridge, M.R. 1998. *Agent Technology Foundations, Applications and Markets*, Springer Verlag.
- Kershenbaum, A. 1993. *Telecommunications Network Design Algorithms*, McGraw-Hill International Editions
- Russell, S. J. and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach*, Prentice Hall, Inc.
- Schoonderwoerd, R. Holland, O.E. and Bruten J. 1996. Ant-like agents for load balancing in telecommunications networks, in *The First International Conference on Autonomous Agents*, ACM Press.
- Stallings, W. 1997. *Data and Computer Communications*, Prentice Hall International, Inc.
- Schwartz, M. and Stern, T.E. 1989 *Routing Protocols, in Computer Network Architectures and Protocols (Second Ed.)*, ed. Sunshine, C.A, Plenum Press, New York / London
- Varian, H.R. 1992. *Microeconomic Analysis (Third Ed.)*, W.W. Norton & Company Inc.
- Varian, H.R. 1995. Mechanism Design for Computerised Agents, *Proceedings of the 1995 Usenix Workshop on Electronic Commerce*.
- M. Wellman, M. 1996. Market Oriented Programming: Some Early Lessons, in *Market-Based Control a Paradigm for Distributed Resource Allocation*, Ed. S. H. Clearwater, World Scientific Press: 74 -95