

Dynamic Resource Allocation by Market-Based Routing in Telecommunications Networks

M.A. Gibney and N.R. Jennings

Department of Electronic Engineering, Queen Mary and Westfield College,
University of London, London E1 4NS, UK.
{M.A.Gibney, N.R.Jennings}@qmw.ac.uk

Abstract. We present an approach to resource allocation in telecommunications networks based on the interaction of self-interested agents which have limited information about their environment. A system architecture is described which allows agents representing various network resources, potentially owned by different real-world enterprises, to coordinate their resource allocation decisions without assuming a priori cooperation. It is argued that such an architecture has the potential to provide a distributed, robust and efficient means of traffic management for telecommunications networks. Some preliminary work on the design of the trading behaviour of the agents in the economy is presented, including the results of experiments which investigate the relative performance of market-based agents compared with traffic management based on static routing.

Keywords: multi-agent systems, telecommunications networks, market-based resource allocation

1 Introduction

In a telecommunications network, a call between two parties may be connected via one of a number of paths. The process of deciding which of these paths to use is called *routing*. Choosing an efficient path is important because the network's capacity for handling traffic is finite, and when it is saturated, calls have to be turned away. This constitutes a loss of income to the network provider. However, finding the optimal path is problematic because the network state continually evolves. By the time the information needed to compute the optimal path between any two nodes is made available at the node where that decision needs to be taken, the network state will probably have changed, rendering that decision obsolete. Furthermore, efficient routing decisions, those which maintain a balance in utilization of the network resources, require information about the utilization of all network resources to be made simultaneously available to the process making that decision.

When a call is requested in a circuit switched network, a call set-up or *connection admission process* tests paths across the network determined by its routing algorithm for congestion and assigns the call to the first path which is uncongested. Routing algorithms are used to establish the appropriate routing paths or the equivalent routing table entries in each node along a path. Most algorithms are based on assigning a cost measure to each link in the network and determining the linear sum of paths across the network [1]. Based on these costs, the network tries to allocate traffic to the cheapest paths across the network. Where the cost function is based on link congestion (real or predicted) the cheapest path may change over time to maintain the network level efficiency of the call admission policy. However, such mechanisms are limited by their lack of information about the wider network state. Put simply, a connection admission process cannot determine the most efficient path from the network point of view merely by checking a small number of paths for congestion. The usage of the links in a path may be equal in two otherwise different network states and the most efficient routing decision from the network point of view would be different in each case. To rectify this situation, we propose a network management framework in which the effects of routing decisions in one part of the network are felt across the entire network, and in which node level decision-making takes place in the presence of some (limited) information about the network level state of congestion. Our approach is to model the telecommunications network resources as trading entities and goods in a computational economy. Although our investigation into this approach is still at the proof of concept stage, we show in this paper that a market-based mechanism is able to successfully route calls through a network over time, although not yet as efficiently as a conventional static routing mechanism. In this paper we discuss the relative performance of the two mechanisms under a variety of network loading conditions and the possible reasons for the under performance of the market-based approach.

One of the main strengths of a market-based approach is that it provides a framework in which the various network components can be owned by different, competitive, real-world enterprises. Thus resource allocation takes place against the assumption of competition, rather than cooperation between the components. - this theme is discussed in section 2. In section 3 we describe the choices that were made in specifying the system architecture and the institutional structure of the market and how these follow from the system level properties we want to achieve. Section 4 describes the design of the individual agents. Section 5 discusses experiments devised to measure the performance of the market-based routing mechanism in a small simulated network under conditions of low, medium and high loading. In section 6 we compare and contrast our approach to that of others who have used market-based mechanisms for network control problems. Finally section 7 describes the open issues and future work.

2 Background and Motivation

A number of agent-based approaches to resource allocation and load balancing in telecommunications networks, for example [2],[3], have worked from the premise that the ideal resource allocation mechanism is one which considers the network as a single resource. In such cases, the system manages the network so as to optimise utilization of that resource. These approaches assume that since global load balancing is a common good, agents should be modelled within a co-operative framework. However, it is not clear that the model of single ownership implied by this co-operative framework will continue to dominate telecommunications deployment in the future. As an economic concern, a network derives its value from connectivity. Thus, large networks grow by acquisition and small networks are swallowed up. However expansion can also happen through co-operative agreement between rival networks to provide inter-connectivity. In the latter case, larger multi-owner networks are created. Another increasingly common aspect of modern telecommunications deployment is the practice of enterprises in other sectors (banking etc.) leasing bandwidth from telecommunications providers. In such an environment, we have the possibility of a number of parties owning resources within a single network which is both horizontally and vertically segmented. Each of these parties clearly has an incentive to see that overuse does not degrade network performance entirely but also an incentive 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 through the setting of policy by sub-network owners within the remit of their own resources. The static nature of these policies and the conflict between them at sub-network interfaces cause institutionalized under-use of the network as a whole.

Given this background, the resource allocation problem in a network with multiple, non co-operating stakeholders can be recast as the problem of reconciling competition between self- interested, information-bounded agents. An effective mechanism for achieving this goal in the real world is the market economy. Therefore in this paper we present some preliminary work towards a telecommunication network management framework modelled on a market economy.

3 System Architecture

This section describes the top-level design of our market-based system. It describes the system architecture (section 3.1), the agent interaction (section 3.2) and the way in which resource commitments are handled (section 3.3).

3.1 A Layered System Architecture

We have designed our system as a three layered model (Figure 1). The lower layer represents a circuit-switched telecommunications network configurable with respect to the number and connectivity of its nodes and links. Resources on this layer are allocated by agents deployed at the nodes in the network which have access to real-time information about the state of utilization of links which are loaded from that node only. The middle layer is the agent-based network management system. It consists of three types of economic agent and two types of market institution which allocate resources in response to the buying and selling behaviour of the economic agents. The third layer is the user layer at which the system interfaces to the call request software. Within the call request software, call prioritization policy can be determined for different priorities of traffic at the user interface by setting spending limits for particular levels of call priority at the source nodes of the network.




Figure 1. System Architecture Layered model of the relationship between the telecommunications network, the market-based multi-agent system and the user.




Figure 2. Network Topology Sample network configuration used for experiments.

In more detail: Link agents sell the underlying resources of the network (i.e. the capacity to carry a traffic along a link). Path agents buy resources from link agents which are then bundled into path resources which can carry traffic across the network. A path agent makes buying decisions within budget constraints to acquire link resources and then subsequently offers to sell the bundles of these resources (as paths) to caller agents which represent the end user. The system has an agent for each link and a market at the source node of that link serving path agents wishing to buy resources on that link. At the path level it has a path agent for each established path across the network (currently three for each source-destination pair) and a path market for each source destination pair in the network. Therefore in the example network shown in Figure 2. there are 24 link agents and link markets, 126 path agents and 42 path markets.

3.2 Agent Interactions

The agents communicate by means of a simple set of signals, which encapsulate offers to sell, bids to buy, commitments of and payments for resources. For convenience, we couple the resources and payments with the offers and bids respectively, this allows for a rapid protocol for resource and information exchange and is used because the time frame in which the resource allocation is carried out is critical to the responsiveness of the system. The speed of the negotiation and allocation protocol is traded off against its lack of flexibility when committing resources. This lack of flexibility may have a detrimental effect on allocative efficiency which in turn could contribute to the under performance of the market-based system across a range of network loadings. Analysis of this trade-off is part of ongoing work.

When a sale is agreed no further communication need take place before the resource can be used or offered for re-sale by the buyer. Buyers and sellers do not communicate directly with one another or amongst themselves. All interactions are by means of institutions (markets). The market institutions also broadcast the clearing prices at which trades are agreed, so the agents have more information upon which to base their trading behaviour. (This has the effect of disseminating network level utilization information to all agents in the system).

The negotiation between buyer and seller agents is mediated by means of a market institution. The type of institution chosen for this purpose in our model was the double auction. A double auction is a market in which buyers and sellers place bids and offers simultaneously, the market matches the bids to the offers thereby setting the price and allocating the goods accordingly. The double auction was chosen for the following properties: It allows multiple buyers and sellers to negotiate simultaneously, it provides a dense set of market clearing price information and it allows supply and demand to be reconciled at the same time [4].

We have chosen to use a variant of the double auction which does not depend on the order or frequency of bids and offers arriving from the participating agents to allow for ultimate migration of the system to a distributed environment in which synchronous message delivery cannot be guaranteed. The variant of the double auction used in the current experiments works as follows: within a trading round, buyers and sellers submit bids and offers for known blocks or allocations of goods. After the auction stops accepting bids and offers (in this experiment agents can submit as many bids and offers as they wish in each time step), it ranks bids from highest to lowest and offers from lowest to highest. Then the highest bid and the lowest offer trade at a price midway between the two. This process continues until no offer remains at a price lower than a bid, at which point all bids and offers remaining are returned to their originators [5]. To make the clearing of resources and payments exchanged in this way efficient, we require that resources be coupled with offers and payments be coupled with bids.

The final role of our market institutions is to ensure acquaintance between agents that are interested in buying and selling particular resources. We distinguish two kinds of markets: one dedicated to mediating the exchange of link resources sold by link agents to competing buyer agents and the other dedicated to selling path resources by path agents to caller agents. These two markets are called the link market and the path market respectively (see Figure 1.).

3.3 Commitment and De-commitment of Resources

Once a price has been agreed between a buyer and seller (through the market institution), the seller's resource is committed to the buyer. The buyer must pay a rental for the use of this resource for as long as it remains allocated to it. In the event that a buyer cannot meet this payment, the resource is de-allocated and reverts to its original owner. A resource which is used profitably (re-sold at a profit) will never be de-allocated under this mechanism - meaning a call cannot be cut-off once it is connected. In addition to de-allocating resources that are not paid for, we also implement a mechanism for the bankruptcy of an agent which is holding resources but not selling them on to callers. Again we prevent calls from being cut-off under this rule by transferring profitable resources to the agent which replaces the bankrupt. Eventually the bankruptcy mechanism will be used to allow the system to learn differential loading patterns and network seasonalities.

4 Agent Design

This section describes the behaviour of the agents, paying particular attention to the functions which determine the prices at which they offer to buy and sell resources. The agents make their bids and offers so as to maximise their income if they are selling and to minimize their expenditure if they are buying. Path agents which both

buy and sell resources, act so as to maximise their profit from ongoing transactions. Gode & Sunder [6] have shown that under certain conditions double auction type markets have the property of generating efficient allocation of resources among zero intelligence agents (i.e. those that do not reason with historical information to predict prices). This result led us to experiment with a relatively simple agent design to determine if routing decisions could be generated by a market mechanism in which agents used pricing functions based solely on endowments and the last known clearing price of their respective markets.

4.1 Link Agents

Link agents can be thought of as the producers of natural resources in the economy. As producers, they wish to maximise the income they derive by selling network resources to path agents. They do this by offering to sell at a relatively low cost when demand is low (in order to continue to acquire contracts) and competition between path agents is less, and at a higher price when demand is high (in order to exploit the greater competition among buyers). The price at which link agents offer resources is set in the following manner:

Let A be the set of link agents, and let $a \in A$ be a specific link agent. Each link agent (a) has a set of resources (R_a) which it can sell in the link market - in this case a unitary resource is sufficient (bandwidth) to carry one call across the link which that agent represents. R_a is composed of n equal bandwidth slices $\{r_{a1}, \dots, r_{an}\}$ which are allocated sequentially from the beginning.

Thus if r_{ap} ($1 \leq p \leq n$) is the next free resource then $p-1$ have already been allocated. After allocating r_{ap} , a will have $n - p$ free resources. a 's pricing function P for the next free resource is given by the following formula:

$$P(r_{ap}) = e^{(p/n)} \quad (1)$$

This function was chosen so that once resources begin to be allocated on a link, the price rises quickly. This means that the marginal cost of obtaining the next resource (when the network is loaded) will be greater than that of obtaining the last.

4.2 Path Agents

Path agents act as both buyers of link resources and sellers of path resources. We detail their buying behaviour in section 4.2.1 and their selling behaviour in section 4.2.2. In general, however, path agents wish to buy resources cheaply from the link agents, and sell them on at a profit to the end-consumers. To do this, they bid competitively to acquire resources, which they then sell on to callers at a price not less than that paid for them. As a secondary factor the agent attempts to sell as many

calls as possible (since these bring in revenue) while keeping inventory at a minimum (since this must be paid for).

4.2.1 Buying Behaviour

A path agent is actually constructed as a buying team which has a centralised budget for the whole path which it divides equally between all team members. Each member of the buying team places a bid to buy resources from a particular link agent at the market at which those link resources are sold.

Let α be a specific path of length n through the network and let PA_α be the path agent responsible for maintaining that path. PA_α will be composed of a team of n sub-agents ($PA_{\alpha,1}, \dots, PA_{\alpha,n}$) each of which is responsible for buying a particular link ($\alpha_i \rightarrow \alpha_{i+1}$) in the overall path.

Assume PA_α has a specific set budget B for buying its needed resources which it divides equally between its team of buying agents. Thus each $PA_{\alpha,i}$ has a budget of B/n . Given the selling price $P_{i \rightarrow i+1}^t$ of a path component at time t , the price bid by $PA_{\alpha,i}$ at the next time instant is given by the following formula:

$$P_{i \rightarrow i+1}^{t+1} = (P_{i \rightarrow i+1}^t + (B / n)) / 2 \quad (2)$$

This function sets the price that the member of the buying team of the path agent offers for a link resources in the centre of the range between the last price paid (lower bound) and the budget for that link (upper bound). More sophisticated strategies are obviously possible.

4.2.2 Selling Behaviour

When a path agent has acquired sufficient resources, from each of its buyer team members, to carry a call across the whole path, it may offer a path resource (i.e. a bundle of an equal amount of resources from each of its buyer team members) for sale to callers. To remain profitable, the agent must not sell these resources at a price any lower than that at which they were acquired jointly by the buyer team. To maximize its profit the path agent should offer to sell this path resource at a price close to (but slightly below) that which it believes callers will be willing to pay. Additionally, the path agent should set its offer price to maximize throughput and minimize inventory.

To be able to offer a path α ($\alpha_1 \rightarrow \alpha_2 \rightarrow \dots \rightarrow \alpha_n$), the path agent has to have bought resources for each of the path links ($\alpha_i \rightarrow \alpha_{i+1}$, for all $1 \leq i \leq n$). Let CO_α be the combined cost paid by PA_α for all the path's constituent components. In putting together the path α , PA_α may have acquired surplus resources which it stores in its inventory IN . PA_α may also have committed a number of path resources to calls

already on the network, we refer to these resources as its throughput TH. The pricing function (see Equation 4. below) is adjusted by a function relating inventory and throughput to the offer price such that the agent offers resources for sale more cheaply when inventory is high and throughput is low.¹

$$f(\text{IN}, \text{TH}) \rightarrow [0,1] \quad (3)$$

Using the price for which a path equivalent to α (i.e. with the same source and destination) was last sold at time t as a guide $P^t(\alpha)$, the price set at the next time instant P^{t+1} ranges between $P^t(\alpha)$ (upper bound) and CO_α (lower bound). (The agent will not offer resources for sale for less than the price for which they were purchased and cannot expect to have a price higher than that set by the auction met in the market). The offer price is set by the following formula:

$$P^{t+1}(\alpha) = CO_\alpha + f(\text{IN}, \text{TH}) \times (P^t(\alpha) - CO_\alpha) \quad (4)$$

At various times, path agents will find that they have bought unequal amounts of the resources they require. In addition to the cost paid to acquire these resources, path agents have to cover the ongoing cost of having these resources allocated to them. To avoid unnecessary costs, a path agent periodically examines its portfolio and returns surplus resources. In this case, surplus means those resources on a given link above the maximum resource allocation which that agent can provide across the whole of its path. More sophisticated mechanisms might have the agent allocate a disproportionate amount of money to acquiring resources to make up a shortfall and thereby balance their portfolio, rather than resort to the wasteful measure of discarding excess, or by making use of the concepts as levelled commitment [7] to allow flexible local deliberation.

4.3 Call Agents

The offers placed on the path market by path agents are allocated to the bids placed on that market by callers. Callers bid for path resources solely on the basis of the most recent price information they have. In future, it may be possible to have callers' bids determined by priority information from an external system (perhaps incorporating the amount of money that a real world caller is willing to pay and / or more sophisticated strategies based on price histories). For now, allowing callers to place their bids near to or slightly above the last price of which they are aware, enables us to successfully route calls across the network.

¹ In our experiments we used a number of functions to generate values in the range 0 to 1 from throughput and inventory state information but did not reach any firm conclusions about their effect on overall system performance.

5 Experiments

To test our hypothesis that a market-based mechanism can effectively balance the traffic loaded onto a network, we conducted a number of experiments, at different network loads, to compare the performance of our system with that obtained using a conventional connection admission control process and static routing. It should be noted that our system does not discover previously unknown routes in the network, rather it seeks to make more efficient use of known routes by incorporating information about the cost to the network of using one route over another. Our simulator consists of a network of seven nodes irregularly connected by twenty-four directed links (see Figure 2). Traffic patterns are described by the total length of simulated time, the mean inter-arrival time between calls, and the mean call duration time. The inter-arrival times and call duration were determined using an exponential distribution. An equal probability of a call request between any source destination pair in the network was used, with approximately one in fifty calls requesting connection between any given source destination pair. We compare the performance of the two systems by measuring the number of calls accepted onto the network in a single time step. Since the total calls offered and mean call holding times are constant in each of our experiments this data can be used to calculate the intensity of traffic carried or lost by each mechanism using the standard telecommunication measure of traffic intensity (the Erlang).

In our simulator we use a negative exponential time distribution function to determine the interval between call arrivals and the duration of calls: Let $U(0, 1)$ be a random distribution function between 0 and 1. The inter-arrival time between calls and the call duration are calculated by the following formula:

$$f(x) = -\beta \ln U \quad (5)$$

Where β is the desired cumulative mean inter-arrival time or call duration, respectively.

The inter-arrival time was varied between 1 and 0.1 seconds the latter representing a dense call arrival. Average call duration was varied between 50 seconds and 500 seconds. Both of which may be considered relatively short call durations. The simulation was allowed to run for 1000 seconds in each case. The interval shown in the graphs (300 - 350s) was chosen because it represents the transition point of loading for the intermediate traffic regime (i.e. the point at which most links in the network become saturated most of the time).

The experimental parameters were chosen to test the trial network using both traffic management mechanisms under a variety of network conditions. We tested the network using traffic ranging from light, in which a low call blocking rate would be

expected (section 5.1), through intermediate (section 5.2) and heavy (section 5.3) in which most calls will not be admitted to the network.

5.1 Sparse call arrival - short call duration

Both mechanisms perform reasonably well under this regime which represents a fairly light load for the network (Figure 3). As can be seen, the market-based mechanism drops more calls on average than the conventional static routing connection admission approach. This is in all probability due to opportunities for trade being missed due to inflexible trading strategies and negotiation protocols. It should be noted that the comparatively good performance of the static routing mechanism under light network loading is due to its sub-optimal decisions having a short duration relative to the inter call arrival time.




Figure 3. Call admission rate of static and market-based routing mechanisms under conditions of sparse call arrival and short average call duration.




Figure 4. Call admission rate of static and market-based routing mechanisms under conditions of intermediate call arrival and intermediate average call duration.

5.2 Intermediate call arrival medium call duration

These results show the static routing network at the transition point beyond which call loss becomes much more likely due to link saturation (Figure 4). The fact that static routing provides a relatively inflexible means of allocating calls to the network means that once network saturation occurs the ability to reroute traffic along alternative routes will not allow the network to make further use of its unused capacity to carry calls.

5.3 Dense call arrival long call duration

The static routing mechanism does not factor in overall performance when allocating a call to a path and therefore the links of the network are allocated to the best paths under light network conditions. Under heavier conditions, these paths lead to further congestion.

At the point where the static routing mechanism is significantly past saturation point the market-based mechanism performs comparably in finding resources on the network to place calls (Figure 5). This is because market-based routing determines the cost of allocating link resources to paths and paths to calls such that it becomes very expensive to use a particular link past a certain point. At that point second and third best routes become cheaper alternatives and this reflects the underlying utilization of the resources throughout the network rather than at one part of it (the path in question).




Figure 5. Call admission rate of static and market-based routing mechanisms under conditions of dense call arrival and longer average call duration.

5.5 Discussion

These results show that the ability of our market-based mechanism to make decisions in the presence of some information about the wider system state is of greatest value in situations of heavy network loading. However, under conditions of light and intermediate loading the market-based mechanism clearly under performed. This may be attributable to the coarse granularity of the resource units relative to the allocation problem being solved, or the fact that the relatively small number of agents meant that the competitive pressure of the market was not as great as it should have been. The simplistic method for de-allocation of resources may also be a contributing factor to the inefficiency of the market-based mechanism described in this work. There is also considerable room for improvement in the design of our agents themselves. The functions which they use to determine the prices to bid and offer for resources are fairly primitive and many opportunities for trade may be lost because of this.

In this work we make some simplifications about the operation of market-based control and static routing that do not properly reflect a real world situation. We allow markets to clear instantaneously and the connection admission decisions of the static routing mechanism to operate instantaneously. This is not sustainable in a real system in the case of static routing because call set up delay plays a significant role in (adversely) affecting the efficiency of the decisions made. It may be possible to overcome this problem in a market-based system by having some of the trading decisions which precede call set up being made before that call is requested through pro-active trading. The effect of time delays on the efficiency of market-based control mechanisms have been studied (particularly in the telecommunication domain) by

Kuwabara, et al. [8] and in the domain of job sharing among networked computers by Chavez, Moukas & Maes [9]. Considerable work remains to be done in assessing the lag time of information passing around the system and the impact of this on allocative efficiency.

6 Related Work

Other authors have applied market-based control mechanisms to network control problems. Most notably Kuwabara, Ishida et al who have investigated an equilibratory approach (in which the market calculates the supply and demand volume equilibrium) to the allocation of resources in an ATM telecommunications network [8]. Our approach differs from theirs in that we ultimately wanted to explore the idea of having the resources needed to make a call across the network ready to be traded at the market at call set up time, thereby avoiding call set up overheads. This led us to treating each call as requiring a unique set of resources and therefore a departure from allocating resources according to a calculated demand-volume equilibrium. Wellman has also described an architecture in which different types of agents internalise the costs of a distributed multi-commodity flow problem by buying and selling resources at prices which reflect the cost to the network of routing traffic through shared links using local information [10]. Wellman's approach yields a balanced flow for the equilibrium state of the network in which the flows across each route are balanced by the market mechanism as predicted by network theory. However, in the domain we wished to investigate (a circuit switched network) each call is discrete and therefore could not be allocated to a number of routes simultaneously (bifurcated routing) as would be required to achieve equilibrium flows across competing routes.

7 Conclusions and Future Work

These preliminary results show the ability of a market-based mechanism to generate paths across a network and allocate calls to them dynamically in an environment in which network resources are owned by a number of different real world entities and therefore co-operation in network management domain cannot be taken for granted. Experimental investigation of the efficiency of this mechanism shows that this technique is most likely to be beneficial when the network is heavily loaded.

Our analysis of the results so far indicates that a more thorough treatment of the relationship between competitive pressures and market efficiency may lead to higher allocative efficiency in the market, and therefore smoother load balancing in the system. Therefore we intend to review how the pricing strategies of the agents are designed and the complexity of the functions required to achieve competitiveness. If the objective were to design the agents for profitability in harsher economic

conditions than those simulated here (in our system callers continue to meet the prices asked by path agents no matter how high they inflate) much could be done to make the path agents more aggressive traders. For example, they could be given the ability to reason about the market over time to adjust their strategies accordingly. An alternative approach which we intend to investigate which makes use of speculating agents to stabilize markets has been shown to provide a worthwhile improvement in allocative efficiency [11].

Another interesting area for future research is the possibility that the market mechanism may allow for the resources needed to make calls to be allocated proactively. This would provide an advantage over conventional call set up / connection admission procedures in that it may be possible to overcome the problem of call set up delay affecting the efficiency of the decision taken. The call set up properties of market-based and conventional routing mechanisms will be investigated in future work.

This work takes place within the wider context of investigation into the application of market-based control to various aspects of telecommunications network management. It is hoped eventually to extend the approach to design a network management system for connection oriented packet switched technologies such as asynchronous transfer mode (ATM) telecommunications networks.

Acknowledgements

This work was carried out under Engineering and Physical Sciences Research Council (EPSRC) grant No.GR/ L04801

References

1. Schwartz M. (1988) Telecommunications Networks - Protocols, Modeling and Analysis, (Addison Wesley. Publishing Company)
2. Appleby S. & Steward S. (1994) Mobile Software Agents for Control in Telecommunications Networks, BT Journal of Technology 12 - 2, pp 104 113
3. Schoonderwoerd, R., Holland, O.E., Bruten, J.L. (1997) Ant-like agents for load balancing in telecommunications networks,. in Proceedings of the First International Conference on Autonomous Agents. Marina Del Ray, California, 1997.ACM Press.
4. Friedman, D. & Rust J. (1991) The Double Auction Market: Institutions, Theories, and Evidence, Proceedings of the Workshop on Double Auction Markets held June, 1991, Santa Fe, New Mexico, Addison-Wesley Publishing Company
5. Davis D. D. & Holt C.A. (1993) Experimental Economics, Princeton University Press
6. Gode D.K. & Sunder S. (1993) Allocative Efficiency of Markets with Zero-Intelligence Traders: Markets as a Partial Substitute for Individual Rationality, Journal of Political Economy Vol 101-1 1993.
7. Sandholm T. (1995) Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework, First International Conference on Multiagent Systems (ICMAS-95), San Fransisco, pp. 328-335.
8. Kuwabara K., Ishida T., Nishibe Y. & Tatsuya S. (1996) An Equilibratory Approach for Distributed Resource Allocation and its Applications to Communication Network Control, in Market - Based Control, A paradigm for distributed resource allocation, Ed. Clearwater S., World Scientific Publishing, 1996, pp 53 - 73
9. Chavez A., Moukas A., Maes P. (1997) Challenger: A Multi-Agent System for Distributed Resource Allocation, in Proceedings of the First International Conference on Autonomous Agents, Marina Del Ray, California, 1997. ACM Press.
10. Wellman M.P, A Market Oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems; Journal of Artificial Intelligence Research 1 (1993) 1 -23.
11. Steiglitz K., Honig M. L. & Cohen L. M. A. (1996) Computational Market Model Based on Individual Action, in Market - Based Control, A paradigm for distributed resource allocation, Ed. Clearwater S., World Scientific Publishing, 1996, pp 1-27