A Microeconomic Approach to Intelligent Resource Sharing in Multiagent Systems

Jaeho Lee and Edmund H. Durfee
Artificial Intelligence Laboratory
The University of Michigan
1101 Beal Avenue
Ann Arbor, MI 48109–2110
{jaeho,durfee}@eecs.umich.edu

Abstract

We have analyzed characteristics of sharable resources and developed techniques for intelligently sharing resources—specifically, communication channels—among agents in multiagent systems. Our techniques allow agents to nearly optimize their communication behavior in a self-organizing and distributed fashion, involving the use of a microeconomic pricing system based on economic laws of supply and demand and trading among agents in real-time. Our analyses are based on three measures of performance: fairness of resource allocation, waiting time for resources, and utilization of resources. Our initial analysis indicates that *fairness* and *utilization* are conflicting, in that the best utilization with a fair allocation is equivalent to the worst utilization with an unfair resource allocation, assuming the allocation policy is statically defined. To strike a balance in performance, we have developed mechanisms that establish an artificial economy, where agents can dynamically reallocate goods (resource access) using a competitive market pricing mechanism. However, unlike more common market-oriented methods, our approach does not demand convergence to equilibrium, but permits more rapid, heuristic trading, leading to near optimal performance where both buyers and sellers of resources can benefit. Our studies show that agents employing our mechanisms can dramatically improve utilization while still providing "fair" access to the resources.

[0]This work was sponsored, in part, by ARPA under contract DAAE-07-92-C-R012.

1 Introduction

In multiagent systems, agents that share an environment can coordinate to solve common problems cooperatively, or to avoid and resolve conflicts that arise as each pursues its separate goals. In this paper, we concentrate on the issue of conflict avoidance, focusing on the use of sharable resources among multiple agents. This type of problem is manifested in many applications involving multiagent systems. For example, agents performing cooperative reconnaissance (Lee, Huber, Durfee, and Kenny 1994) must share their communication resources so that the most important messages are exchanged without interference or collision. Similarly, mobile robots that use similar active sensors (such as sonar) could be confused by crosstalk (sensing others' signals) unless they somehow avoid the simultaneous use of their sensing resources (Kortenkamp, Huber, Koss, Lee, Wu, Belding, Bidlack, and Rodgers 1994). In essence, then, the problem we must address is how agents might make intelligent decisions about the use of shared resources such that each can have a better chance of using the resources when they are really needed.

In this paper, we present techniques for intelligently sharing resources—specifically, communication channels—among agents in multiagent systems. Our techniques allow agents to nearly optimize their communication behavior in a self-organizing and distributed fashion, involving the use of a microeconomic pricing system based on economic laws of supply and demand and trading among agents. The agents, therefore, allocate resources based on an artificial economy or market (Wellman 1993; Bogan 1994; Kuwabara and Ishida 1992). In work following a similar microeconomic model, Wellman (1993) states that market-oriented programming refers to the general approach of deriving solutions to distributed resource allocation problems by computing the competitive equilibrium of an artificial economy. This approach is distinguished from general search-based approaches such as Sycara, Roth, Sadeh, and Fox (1990), where agents use sophisticated local control. In fact, the notion of using microeconomic methods to allocate communication resources dates back at least to the work of Kurose, Schwartz, and Yemini (1985), which described how a pricing system and a perfectly competitive market could be used to compute, in a distributed fashion, an optimal solution to the resource allocation problem.

Kurose's techniques, and more generally the market-oriented programming approach, rely on the computation of a market equilibrium to establish resource allocation policies. While potentially yielding optimal allocation decisions, equilibration can at times require many iterations to converge, if it ever does so. If, in the meantime, resource requirements change, or in cases where resource access must be nearly instantaneous to be useful, waiting for equilibrium might be inappropriate. For example, the utility sending a message might be very time-dependent, so the communication decision should be made rapidly. More generally, the communication resource needs of an agent will vary over time, as it can experience periods of relative quiet, punctuated by bursts of communication. For such cases, agents in a multiagent system could use less expensive, heuristic algorithms that guarantee rapid, near-optimal resource allocation based on the constantly evolving resource needs of the agents. Our work, as described in this paper, provides a principled set of heuristics that agents can use in dynamically buying and selling resources among themselves. With regard to communication resources, this amounts to agents "buying" the silence of others, but those others can in turn accumulate enough of their own buying power to commandeer resources in exceptional circumstances. We formalize and operationalize this notion in the remainder of this paper.

2 Communication Resource Allocation

While the approach we develop is applicable to a range of domains that involve the sharing of resources such that no agent benefits if multiple agents attempt to access the resource simultaneously, the particular domain that we will focus on is communication resource allocation. More specifically, the domain we are exploring involves *port-based communications* and $p-$persistent CSMA protocols (Tanenbaum 1981) with multiple channels with different communication delays. In this domain, agent$_a$ has ports 1 to N connected to the corresponding ports of agent$_b$. When an agent becomes ready to send, it senses the channel. If the channel is idle, the agent transmits with probability p. With a probability $q(= 1 - p)$ it defers until the next slot. If that slot is also idle, it either transmits or defers again, with probability p and q, respectively. This process is repeated until either the message has been transmitted or another agent has begun transmitting. For our analyses, we further assume that: (1) all messages are of constant length; (2) there are no errors except those caused by collisions; (3) there is no capture effect (Tanenbaum 1981), so a collision leads to the loss of all colliding messages; and (4) the state of the channel can be sensed instantaneously.

2.1 The Need to Trade Resource Access

At first, let us consider only *one* channel shared by two agents, agent$_a$ and agent$_b$. Assume that, when the channel is idle, each agent attempts to send a message with transmission-potential P_a and P_b, respectively. If each agent has an *infinite* number of messages in the queue, the utilization of the channel, ρ, is thus given by

$$\rho = P_a(1 - P_b) + P_b(1 - P_a) \tag{1}$$

If the two agents have the same transmission-potential ($P_a = P_b$), Equation 1 becomes $2P(1 - P)$, where $P = P_a = P_b$. This relation is shown in Figure 1. The *maximum* channel utilization occurs at $P = 0.5$, with $\rho = 0.5$. That is, half of the time exactly one agent is transmitting. The other half of the time either both or neither is transmitting, and no communication occurs.

On the other hand, if the sum of transmission-potentials is fixed to a constant, say $P_a = 1 - P_b$, Equation 1 becomes $2P^2 - 2P + 1$, where $P = P_a = 1 - P_b$. The shape of the curve is shown in Figure fig:utilization-b. The *minimum* utilization occurs at $P = 0.5$, with $\rho = 0.5$, and maximum

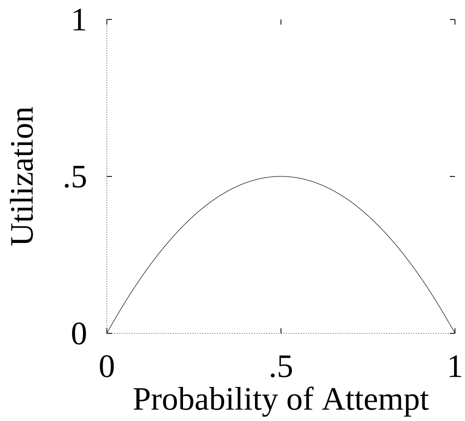

Figure 1: $P_a = P_b$

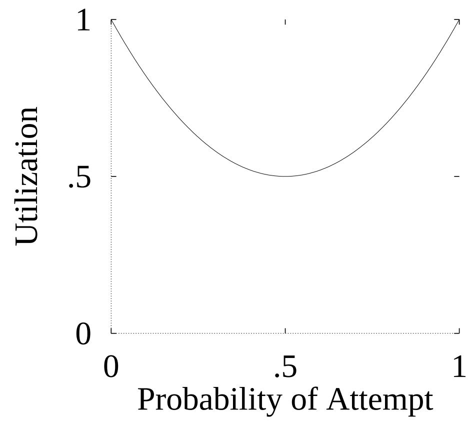

Figure 2: $P_a + P_b = 1$

channel utilization can be achieved by assigning all potential to one agent. That is, if one has all the potential, it is always transmitting while the other is always quiet, and the channel is always being used successfully.

In summary, when we assign equal transmission-potentials to the two agents, *maximum* throughput is achieved by giving them 0.5 transmission-potentials. When the sum of transmission-potentials is fixed to 1, and agents can have unequal portions, the total throughput becomes *minimal* at 0.5 by making the transmission-potentials equal to 0.5 for each agent. From this analysis, we can conclude that when there are infinite messages to send, the best "fair" allocation of potentials corresponds to the worst "unfair" allocation.

Throughput, however, is not the only parameter we are interested in. The mean message delay is also important, especially for messages with deadlines. Unfortunately, high throughput and low delay are inherently in conflict. Good performance on one of them can be achieved only at the expense of the other. For example, if one agent has all the potential and the other has none, the utilization of the channel is maximized to 1 (assuming agents have infinite messages to transmit), but the mean message delay becomes infinity. The simplest way to overcome this problem is to alternate transmission potentials between the two agents. In other words, agent$_a$ has a transmission-potential of 1 and agent$_b$ has 0 potential for a while, and then they switch their potentials.

2.2 Trading Resource Access

Simply swapping 1/0 potentials works very well with infinite messages, but what if the number of messages to send is not infinity? At this point, we relax the assumption that agents have infinite messages, and instead assume that the probability that an agent has a message to send is M_a for agent$_a$ and M_b for agent$_b$. In this case, some bandwidth of the channel may be wasted because the agent may have no message to send while it is holding all the potential. Rather than having some static "turn-taking" behavior, the system would be more adaptive if agents could "trade" potential as needed.

Our approach is thus to have an agent *trade* surplus transmission-potential if another agent is willing to purchase more potential. In our model, each agent has a *initial budget*, and a regular *income*. The agent pays the channel for delivering messages, but it does not pay when the messages collide. The *price* of the channel is changed periodically by economic laws of supply and demand. Since the budget is limited by the initial budget and the regular income, and each agent has to pay the channel to send messages, one agent cannot hold the entire transmission-potential all the time.

In the following sections, we investigate issues involved with our trading model. Section 3 examines the properties of transmission-potential in detail. Section 4 investigates the pricing mechanism of the channels. Section 5 deals with the issues related with the income of the agents. Section 6 explains the algorithm for trading transmission-potential. Section 7 describes our preliminary and planned evaluation and Section 8 outlines the current status of this work and our ongoing research.

3 Transmission Potential

In Section 2.1, we analyzed the simple communication model on the assumption that each agent has infinite messages to send. Now suppose that the probability that an agent has a message to send is M_a for agent$_a$ and M_b for agent$_b$. The utilization is now expressed in terms of M_a and M_b:

$$\rho = M_a P_a (1 - M_b P_b) + M_b P_b (1 - M_a P_a) \tag{2}$$

Figure 3, 4, 5 show the contour lines of utilization, ρ, versus P_a and P_b for simple cases where $M = M_a = M_b$. Figure 3 shows that for $M \leq 0.5$ the best utilization is achieved by keeping both of the agents' transmission-potential at 1.0. In other words, $P_a + P_b$ can be 2.0 for $M \leq 0.5$. For $M > 0.5$, the contour lines in Figure 4 and 5 show that there are two separate regions with $\rho \geq 0.5$; that is, an upper left region and a lower right region. Let β be the boundary value of P_a and P_b of these regions, $\beta = 1/(2M)$. The upper left region is bounded by $P_a \leq \beta$ and $P_b \geq \beta$ and the lower right region is bounded by $P_a \geq \beta$ and $P_b \leq \beta$. Let us define these regions as *goal-regions*. Once again, from the location of goal-regions, we can see that the more we bias the transmission-potential, the higher utilization we can achieve. Note that the goal-regions are convex and thus have ridges. In other words, the utilization at the ridge of the region is relatively higher than the rest of the region. We can approximate these ridges using a curve[1]

$$P_a = (1 - P_b^{-\log_\beta 2})^{-\log_2 \beta}, (0.5 \leq \beta < 1.0) \tag{3}$$

[1] The curve can expressed by $P_a^S + P_b^S = 1$, and this equation has to pass (β, β); that is, $\beta^S + \beta^S = 1$. This leads to $S = -\log_\beta 2$.

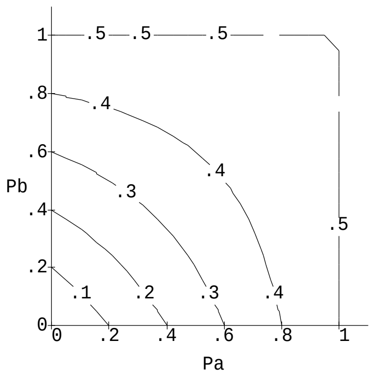

Figure 3: $M = 0.5$

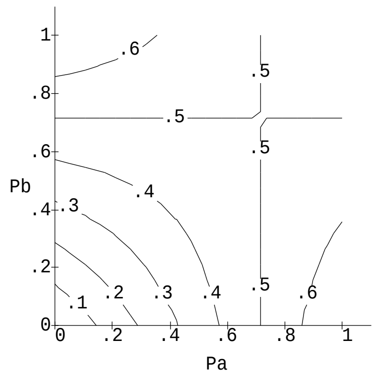

Figure 4: $M = 0.7$

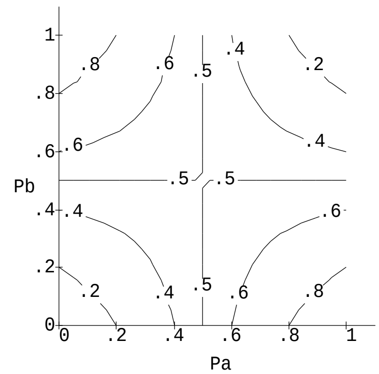

Figure 5: $M = 1.0$

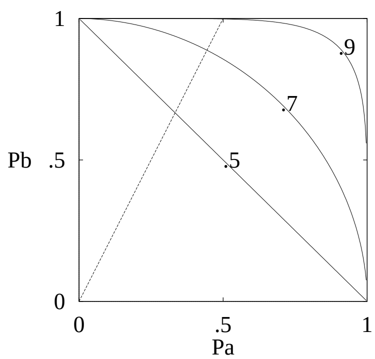

Figure 6: Potential Ridges

$$\gamma = \frac{P_b^t}{P_a^t}$$

$$P_a^{t+1} = (1 + \gamma^{-\log_\beta 2})^{-\log_2 \beta} \quad (4)$$

$$P_b^{t+1} = \gamma(1 + \gamma^{-\log_\beta 2})^{-\log_2 \beta} \quad (5)$$

Figure 7: Calculations of New Potential

which passes three points (0,1), (β, β) and (1,0), and lies in the *goal-regions*. Let us call this curve the *potential-ridge*. Figure 6(a) shows potential-ridges for $\beta = 0.5, 0.7$, and 0.9 from left to right.

From these graphs, a simple *heuristic* for computing the total transmission-potential of a port, T, can be derived. Let us suppose that the agents agree that the overall utilization of channels should be greater than 0.5. If the probability that the agents have messages to send is greater than 0.5, then T can be as low as 1.0[2], and if the probability is less than 0.5, T can be as high as 2.0 since each agent can hold 1.0 of transmission-potential. In other words, we can say $1.0 \leq T = P_a + P_b \leq 2.0$ and T is a function of β.

Since in real networks the value of M is not readily available, the value of β is not available, either. Nevertheless, we can measure the past demand for the channel to approximate β. The total demand for a channel is increased whenever agents attempt to use the channel. Hence, if two agents

[2]A fair strategy might assign 0.5 potential to each agent, and a biased strategy might assign 1.0 and 0.0 potential to each agent, respectively.

try to send messages at once (i.e. collision), the demand for the channel is increased by 2, and if only one agent attempts to use the channel the demand is increased by 1. The *demand-ratio* of a channel, D, is the number of demands per unit time. From the *heuristic* for the value of T, the demand-ratio (0,2) can be mapped to T (2,1) by the following simple transformation equation, $T = 2 - D/2$. Once the value of T is derived, we have a potential-ridge corresponding to the new T. Since $T = 2\beta$, we can derive the estimated value of β as $\beta = T/2$.

The increase or decrease in T is shared by the agents. The ratio of sharing is determined by the current P_a, P_b and the new potential-ridge. The new transmission-potentials of the agents, P_a^{t+1} and P_b^{t+1} at time $t+1$ are the contact point of the new potential-ridge and the line which connects the origin (0,0) and the point of previous transmission-potentials (P_a^t, P_b^t). This relation is shown in Figure 6, and the equations for the new potentials in Figure 7. These equations guarantee that the new transmission-potentials are on the potential-ridge, and hence in the goal-regions.[3]

4 Prices of Channels

Until now, we have analyzed the model where two agents share only one communication channel. Now let us return to the our original model of multiple channels. There are N different channels and each channel has a different service-rate and price. Whenever an agent sends a message, the agent has to pay the chosen channel for delivering the message. The payment is determined by the price of the channel. The price of a channel rises as the demand for that channel increases, and falls as the demand decreases. The price, R_i^t, of the channel$_i$ connected to the port$_i$ at time t can be defined as the *demand-ratio* D_i^t of channel$_i$ at time t (Section 3), $R_i^t = D_i^t$. Recall that the *demand-ratio* has value in the range of (0,2). This equation exactly represents the economic laws of supply and demand.[4]

Changing prices of channels has two important effects. First, in our model where multiple channels are shared by two agents, the pricing mechanism affects load-balancing between multiple channels because the agents try to select a channel which can send a message at the lowest price within a deadline. Second, if the prices of all channels go up as the demand for channels increases, only agents who can afford the increased prices can send messages. This is a very important feature of the pricing mechanism, and it will be discussed more in Section 6.

[3] Note that simply increasing or decreasing transmission-potentials at the ratio of current transmission-potentials can cause one agent to have more than 1.0 potential, which does not make sense.

[4] We may introduce a *sensitivity-factor* σ to make $R_i^t = \sigma D_i^t$ for some real constant σ. In this case, the price has value between 0 to 2σ.

5 Income of the Agents

In the previous sections, we have seen that the best utilization of channels can be achieved by a biased allocation of transmission-potentials between agents, but a biased allocation can cause indefinite transmission delays. Our approach to solving this problem is to give each agent a regular "income" over time and to have an agent spend from its budget for transmitting messages. Whenever an agent transmits a message, the agent has to pay the channel, and the payment is determined by the price of the channel. The agents are allowed to sell their surplus transmission-potential to other agents who need more potential, in return for some of the income of those agents. By selling potential, an agent enlarges its budget so that it can buy a lot of potential when it needs to in the future.

In our distributed communication model, the role of the budget is very important. While the channel prices balance the loads on the channels, the budgets of the agents promote trading potentials and sharing channels. Another important issue is to maintain a proper budget level for each agent. If both agents always have enough budget to send all messages and the income is also sufficient, then trading transmission-potential may not occur. Maximal benefit of trading potentials to properly utilize the channels only comes about if the total budget of the agents is just enough to use the capabilities of the channels.

One way of avoiding excessive budget is adjusting the income of the agents on the basis of the "channels' income" which is the same as the budget paid by the agents. The model we have employed is a kind of "zero-sum" society where the total worth of the system is basically preserved. Let us specify the system parameters of our model in detail before the equations for the new budget and the income of the agents are presented.

- Trading occurs periodically, and the length of the period is L. In our notation, superscript t or $t+1$ means the current period and the next period, respectively.
- The system has N ports from port$_1$ to port$_N$ for each agent and each port is connected to N channels from channel$_1$ to channel$_N$. channel$_i$ has S_i message service-rate (messages/time). channel$_i$'s price at time t is R_i^t.
- *Utilization goal*, $0.0 \leq U_g \leq 1.0$ is defined by the system and specifies the system-wide maximum utilization goal of all channels. Recall that higher U_g values lead to longer average message delays.
- *Virtual service-rate*, $V_s = \sum_{i=1}^{N} S_i$ is the total service-rate of all channels.
- *Virtual service-price*, $V_p^t = \sum_{i=1}^{N} S_i R_i^t / V_s$ is the average price per message over all channels (price/messages).

- Each agent has initial budget, B^0, before it enters the system, $B^0 = U_g \times L \times V_s \times V_p^0/2$ where V_p^0 is the initial *virtual service-price* of the channels which is the function of the initial channel prices.

- *Income of the channels*, I_{ch}^t is the sum of the budget paid by both agents for the previous period.

- *Income of an agent* for the next period, $I_{ag}^{t+1} = I_{ch}^t/2$ is the dividend of the channels' income for the previous period.

Note that agents can calculate all of these values without any communication with the other agent. For example, the income of the channels, I_{ch}^t and their prices, R_i^t, can be obtained just by monitoring the busy status of channels.

6 Trading Transmission Potential

A trade of transmission-potential occurs only when trading benefits both agents and they agree to do so. Agents communicate periodically with all the necessary information to trade potentials.[5] The necessary information is called BCM. B represents the *B*udget available for the next period, which has been derived in the previous section. C represents the transmission *C*apacity calculated in budget units, and M represents the budget required to send all expected *M*essages for the next period. In other words, B is how much an agent can spend, C is how much an agent would spend to use its full potential, and M estimates how much an agent needs to spend to send all the messages. Note that we have transformed the BCM information to the same "budget" units. To derive the equations for C and M, we need to define more system parameters as follows:

- An agent has P_i^t transmission-potential for port$_i$ at time t.

- *Estimated message-generation rate*, $G_e^t = hG_e^{t-1} + (1-h)G_r^{t-1}$ of an agent is a weighted sum of past and current message-generation rates. The message generation rate, G_r^{t-1} is measured by counting the number of messages generated during the previous period. The constant h, for example 1/2, determines how quickly the agent reacts to recent history.

- *Expenditure for a port$_i$*, $E_i^t = S_i P_i^t R_i^t$ is the expected budget of sending infinite number of messages with the current transmission-potential at the current price of channel$_i$.

[5] Agents send a message for trading to each other, and the deal occurs at both sides at once. Since the deal is deterministic, agents need not communicate again.

Case	Action	Amount	
$(1) C, M \geq B$	Sell potential	$(C-B)/2$,	if $(C-M) \leq (M-B)$
		$(C-M)$,	if $(C-M) > (M-B)$
$(2) B, M \geq C$	Buy potential	$(B-C)/2$,	if $(B-M) \leq (M-C)$
		$(M-C)$,	if $(B-M) > (M-C)$
$(3) B, C \geq M$	Sell potential	$(C-M)$	

Table 1: Heuristic Criteria for Trading Decision

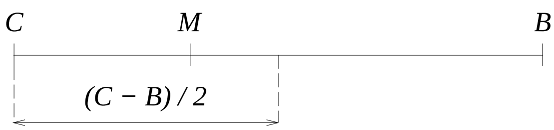

Figure 8: $(C-M) \leq (M-B)$

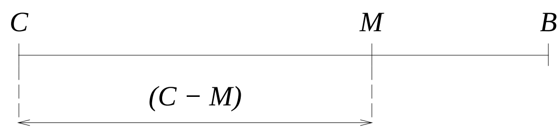

Figure 9: $(C-M) > (M-B)$

Now BCM information can be derived as follows:

$$B^{t+1} = B^t + I_{ag}^{t+1} \qquad C^{t+1} = L \times \sum_{i=1}^{N} E_i^t$$

$$M^{t+1} = V_p^t \times (L \times G_e^t + \text{\# of messages in the queue})$$

Note that each agent calculates its BCM values independently, and the agents exchange this information to decide on trades. Table 1 shows heuristic criteria for trading decisions of agents based on the BCM information. A trade occurs only when one agent wants to sell transmission-potential and another agent is willing to buy it. If the amount of the trading offer of each agent is different, the agents trade only the *minimum* amount of their trading offers. As mentioned earlier, trading can benefit both agents. For example, if an agent has $C > M > B$, the agent has sufficient transmission-potential and channel capacity C to send all messages M, but its budget B is insufficient to use all of its capacity. Hence the agent should trade away its some of its unusable capacity to get more budget. If $(C-M) \leq (M-B)$ (Figure 8), the best deal for the agent is selling $(C-B)/2$, thereby increasing its budget by that much. If the deal is made, the agent uses its larger budget to send more. If $(C-M) > (M-B)$ (Figure 9), the agent can improve its throughput by $(M-B)$ by selling all surplus transmission-potential, $(C-M)$. Other cases in Table 1 can be explained similarly. Especially, case (3) is interesting. In case (3), the agent has enough budget and capacity to send all expected messages, but it tries to sell its extra potential to increase its budget for future use. Consider a scenario where each agent has the same amount of budget. Initially, agent$_a$ generates many messages to send and agent$_b$ does not. By the trading heuristic given in the Table 1, agent$_b$ sells some of its potential to increase its budget. Later, agent$_a$ is still generating many messages but now agent$_b$ also starts to generate many messages. Since both agents are more likely to be sending messages at the same time, contentions for the channel drives up its price. Now since their incomes are limited, only agents who have accumulated higher budgets (such as agent_b) can afford the raised prices, while agents that have been spending their budgets (such as agent_a) have to sell off transmission-potential in order to build up their budgets. This is the reason that each agent tries to keep as high a budget as possible. Once the amount of trading is determined, each agent increases or decreases its transmission-potential for each port in the ratio of the current transmission-potential assignment.

In summary, each agent calculates the BCM values using only local information. Each agent then attempts to maximize its throughput by trading with other agents. An effective global throughput of the system is thus achieved indirectly through the selfish behavior of each agent.

7 Evaluation Summary

A simulation program for testing the effectiveness of our approach has been built using the CLOS (COMMON LISP Object System). Using the simulation program, we have been testing the correctness of our algorithms and analyzing the pricing and trading heuristics. We are also planning to conduct comparative experiments with other knowledge-poor, systematic approaches such as the simple turn-taking method and the *binary exponential backoff* protocol heuristic (Tanenbaum 1981). Such comparative experiments have proven problematic because of the many communication parameters that need to be specified, including message generation characteristics, message utilities and deadlines, channel service rates, etc.

We have focused on measuring performance along the dimensions of fairness, resource utilization, and agent utility from the resource (factoring in message delivery delay). Our preliminary evaluations have highlighted how, with fair incomes among agents, agents that are more taciturn receive higher utility when they do communicate; that selling their usual silence leaves them with a high budget to get messages through when they need to. A more talkative agent, on the other hand, has a greater access to the communication channels most of the time. This has been contrasted to either of the two “fair” static approaches described earlier: either each agent holds an equal fraction of the channel potential (so that the one that has more to say does not say it as frequently, and channel utilization is poor); or agents flip-flop between them (for equal amounts of time), such that, once again, one has more resource than it needs and the other cannot make use of the excess resource. Our ongoing work is in using our simulation results to quantify performance improvements in well-characterized scenarios, and incorporate measures of trading overhead into our assessments.

8 Conclusion

In this paper, we have examined an approach towards intelligently sharing communication resources in multiagent systems based on models and methods from the field of economics. We showed how to map the resource sharing problem into artificial economy model, defined what to trade and how to trade, and experimentally shown how trading between agents allows a beneficial solution for multiple parties.

Our model and algorithm is relatively cheap and simple, and thus applicable for dynamic real-time communications where other market-oriented mechanisms requiring equilibration among agents would be infeasible. Because coordination over resource use occurs dynamically, run-time communication overhead must be considered, but in our algorithm the number of communication slots used for trading transmission-potentials is only one per trade, and other necessary information for adapting to new communication situations is obtained by the usual monitoring of channels.

Our preliminary analyses and the model presented in this paper need many extensions and much evaluation to be used in real multiagent systems. First of all, more complete experimentation is needed. Second, our analysis here is based on interactions between two agents, for simplicity. With more agents, the trading heuristics will become more complex, but the underlying philosophy will continue applying. Third, in highly competitive settings, protocols for revealing less information, or that enforce truthful revelation (Rosenschein and Zlotkin 1994), must augment or replace the BCM information exchange presented here.

References

Bogan, N. R. (1994, May). Economic allocation of computation time with computation markets. Master's thesis, Massachusetts Institute of Technology.

Kortenkamp, D., M. Huber, F. Koss, J. Lee, A. Wu, W. Belding, C. Bidlack, and S. Rodgers (1994, March). Mobile robot exploration and navigation of indoor spaces using sonar and vision. In *Conference on Intelligent Robotics in Field, Factory, Service, and Space (CIRFFSS '94)*, Houston, Texas, pp. 509–519.

Kurose, J. F., M. Schwartz, and Y. Yemini (1985, May). A microeconomic approach to optimization of channel access policies in multiaccess networks. In *Proceedings of the Fifth International Symposium on Distributed Computing Systems*, pp. 70–80.

Kuwabara, K. and T. Ishida (1992, July). Symbiotic approach to distributed resource allocation: Toward coordinated balancing. In A. Cesta, R. Conte, and M. Miceli (Eds.), *Pre-Proceedings of the Fourth European Workshop on Modeling Autonomous Agents in a Multi-Agent World*, Rome, Italy.

Lee, J., M. J. Huber, E. H. Durfee, and P. G. Kenny (1994, March). UM-PRS: an implementation of the procedural reasoning system for multirobot applications. In *Conference on Intelligent Robotics in Field, Factory, Service, and Space (CIRFFSS '94)*, Houston, Texas, pp. 842–849.

Rosenschein, J. S. and G. Zlotkin (1994, Fall). Consenting agents: Designing conventions for automated negotiation. *AI Magazine 15*(3), 29–46.

Sycara, K., S. Roth, N. Sadeh, and M. Fox (1990, November). Managing resource allocation in multi-agent time-constrained domains. In *Proceedings of the 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, San Diego, CA, pp. 240–250.

Tanenbaum, A. S. (1981). *Computer Networks.* Prentice-Hall.

Wellman, M. P. (1993). A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research 1*, 1–23.