Characterizing Contract-Based Multiagent Resource Allocation in Networks

Bo An, *Student Member, IEEE*, and Victor Lesser, *Fellow, IEEE*

***Abstract*—We consider a multiagent resource allocation problem where individual users intend to route traffic by requesting the help of entities across a network, and a cost is incurred at each network node that depends on the amount of traffic to be routed. We propose to study contract-based network resource allocation. In our model, users and nodes in the network make contracts before nodes route traffic for the users. The problem is an interesting self-interested negotiation problem because it requires the complete assembly of a set of distinct resources, and there are multiple combinations of distinct resources that could satisfy the goal of negotiation. First, we characterize the network allocation problem and show that finding optimal allocations is $\mathcal{NP}$-complete and is inapproximable. We take both Nash equilibrium and pairwise Nash equilibrium as the solution concepts to characterize the equilibrium allocations. We find that, for any resource allocation game, Nash equilibrium and pairwise Nash equilibrium always exist. In addition, socially optimal allocations are always supported by Nash equilibrium and pairwise Nash equilibrium. We introduce best-response dynamics in which each agent takes a myopic best-response strategy and interacts with each other to dynamically form contracts. We analyze the convergence of the dynamics in some special cases. We also experimentally study the convergence rate of the dynamics and how efficient the evolved allocation is as compared with the optimal allocation in a variety of environments.**

***Index Terms*—Negotiation, networks, resource allocation.**

I. Introduction

In SYSTEMS involving multiple autonomous agents, it is often necessary to decide how scarce resources should be allocated. The allocation of resources within a system of autonomous agents is an exciting area of research at the interface of computer science and economics. Multiagent resource allocation (MARA) is relevant to a wide range of applications, e.g., industrial procurement, manufacturing systems, and grid computing (see [4] for a survey). Market mechanisms have been evolving as a method to incorporate control into systems with a large number of autonomous agents. Economic paradigms allow for decentralized implementation while providing mechanisms to regulate the behaviors of users. There have been extensive applications and support for market-based control, e.g., factory scheduling [5] and communication networks [16].

This paper looks at a MARA problem across computational networks consisting of selfish entities. Individual users send traffic by requesting the help of those entities. A cost is incurred at each node (entity) receiving, storing, and processing traffic. In this paper, we study contract-based mechanisms for the resource allocation problem. In our model, users and nodes in the network make agreements (contracts) before nodes route traffic for the users. The model we develop is motivated by multiagent environments, where bilateral negotiations may result in service agreements between users and network entities (nodes). Agreements are used for allocating traffic flow, and network entities' costs are compensated in the form of payments associated with agreements. Our problem is an interesting self-interested negotiation problem because it requires the complete assembly of a set of distinct resources, and there are multiple combinations of distinct resources (i.e., multiple paths) that could satisfy the goal of negotiation.

The model considered in this paper can be motivated from various contexts. As an example, consider a distributed vehicle network that consists of a number of geographically dispersed dispatch centers of different companies [28]. Each center is responsible for certain deliveries and has a certain number of vehicles. There is a cost associated with each center, which represents the storage and labor costs of handling the delivery tasks. When a user wants to deliver certain traffic through a path, it needs to make contracts with all delivery centers on that path. When a user has a delivery task, there may be multiple paths from the source to the destination. The user just needs to request the centers along any path to finish the delivery task.

As another example, consider the negotiation management component [1] for Collaborating, Autonomous Stream Processing Systems (CLASP) [3], which has been designed and prototyped in the context of System S project [14] within IBM Research to enable sophisticated stream processing. There are multiple sites running the System S software, each with their own administration and goals. Each site may only have limited processing capabilities, so cooperation among these sites can frequently be of mutual benefit. Consider that a site receives a job. After planning [22], the site finds that using only its local resources, it cannot satisfy all resource requirements of the plan. Then, the site negotiates with other sites to acquire resources that are needed, and such negotiations are conducted by the site's negotiation management component [1]. The plan can be executed if and only if *all* resource requirements are satisfied. There may exist multiple functionally equivalent

Manuscript received December 20, 2008; revised June 16, 2009. First published December 1, 2009; current version published June 16, 2010. This work was supported primarily by the Engineering Research Centers Program of the National Science Foundation (NSF) under NSF Cooperative Agreement EEC-0313747. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the NSF. This paper was recommended by Associate Editor T. Vasilakos.

The authors are with the Department of Computer Science, University of Massachusetts, Amherst, MA 01003 USA (e-mail: ban@cs.umass.edu; lesser@cs.umass.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSMCB.2009.2035100

plans, and different plans require different set of resources. To complete the job, the site needs to acquire the resources of any plan by negotiation. There is a cost associated with a site's providing resources, and the cost depends on factors like the number of sites using its resource. Each plan in this stream processing example can be treated as a path in our problem.

Our first objective in this paper is to understand the relationship between stability and optimality in this resource allocation game, where selfish agents strategically choose their actions. While selfish agents are only interested in maximizing their individual objective functions, the overall efficiency of *stable* allocations that are outcomes of decentralized strategic interactions can be worse than that of social *optimal* allocations formed by a central authority maximizing aggregate social welfare. We show that finding optimal resource allocations is $\mathcal{NP}$-complete and is inapproximable. We analyze both Nash equilibrium (NE) and pairwise NE as the solution concepts to characterize the equilibrium allocations. While NE is the dominant solution concept in studying agents' strategic behavior in various resource allocation games, we also study pairwise NE. This is a refinement of the NE, which requires that an NE stable allocation is immune to the strategy deviation of any pair of agents [12]. We find that, for any resource allocation game, NE and pairwise NE always exist. In addition, the optimal allocations are always supported by NE and pairwise NE.

The other main objective of this paper is to understand the dynamics of the resource allocation process when a collection of agents interact with each other to dynamically form contracts. We explore how the allocations evolve under such dynamics in light of the underlying resource allocation game. We are interested in understanding and characterizing the allocations that result when the decision-makers interact to choose their actions. We define protocols in which agents can dynamically respond to current allocation by taking myopic best-response strategies. Individuals do not predict how their contracting decisions might affect future decisions of other individuals or, more generally, how they might influence the future evolution of the allocation. Such myopic behavior occurs in large multiagent systems where agents have limited information about the incentives of others or where agents are resource bounded. In this paper, we analyze the convergence property of the dynamics in some special cases. In addition, we consider decommitment penalty [29] as an additional variable for an agent to decide the set of allowable contracting actions. We also experimentally study the convergence rate of the dynamics and how efficient the evolved allocation is as compared with the optimal allocation in a variety of environments.

In our approach, users and network nodes make proposals and then reach agreements. Our negotiation approach for MARA is of a *distributed* nature. In general, the allocation procedure used to find a suitable allocation of resources could be either *centralized* or *distributed*. The typical and best-known examples for the centralized approach are combinatorial auctions [6], [20]. In combinatorial auction schemes, a centralized controlling agent (the "auctioneer") assumes responsibility for determining which agents receive which resources based on the bids submitted by individual agents. However, the problem of deciding successful bids, i.e., winner determination problem, is $\mathcal{NP}$-hard [24]. In addition, the auctioneer may face significant computational overload due to a large number of bids with complex structures. One of the most important arguments against centralized approaches is that it may be difficult to find an agent that could assume the role of an "auctioneer." For instance, selfish resource providers may not trust the auctioneer and are not willing to comply with the decisions made by the auctioneer. In distributed approaches like automated negotiation, on the other hand, allocations emerge as the result of a sequence of distributed negotiations, and each selfish agent acts on behalf of itself. The distributed model seems more natural in cases where resources belong to different selfish agents, and finding optimal allocations may be (computationally) infeasible.

Although users and network nodes "cooperate" with each other on the form of contracts, the focus of this paper is not to investigate the agents' strategic behaviors under complex bargaining protocols (e.g., infinite time alternating offers bargaining [26, p.100]). Instead, we are trying to characterize the network resource allocation game based on using straightforward contracting protocols. While analyzing agents' equilibrium strategies, agents simultaneously announce their bids (prices in this paper). While studying the best-response dynamics, an agent makes a proposal first, and the agent accepting the proposal either accepts or rejects the proposal. In addition, each agent is allowed to decommit from existing agreements. While modeling resource allocation dynamics, each agent adopts a best-response myopic strategy. In game theory, the best response is the strategy (or strategies) that produces the most favorable outcome for a player, taking other players' strategies as given. In evolutionary game theory, best-response dynamics represents a class of strategy updating rules, where players' strategies in the next round are determined by their best responses to some subset of the population. Myopic best response refers to dynamical rules in which players do not consider the effect that choosing a strategy on the next round would have on future play in the game. Myopic best-response strategy is widely used to model selfish agents' natural behaviors, particularly bounded rational agents [21], and, in this paper, we will use it to model the formation of stable allocations in a distributed way.

Like most formal papers on MARA (e.g., [8] and [9]), selfish routing (e.g., [25]), and network formation (e.g., [12], [15], and [31]), we restrict our attention to games with complete information where each agent has complete information about the other agents. Although this assumption may seem impractical, our problem is already complex enough given agents' selfishness and strategic interaction. Moreover, our analysis will provide us insights into more general resource allocation problems, e.g., factors affecting the resource allocation results. In our experiments, we show how some simple forms of uncertainty affect the resource allocation result. Our result is a good starting point for studying more realistic models with incomplete information.

The remainder of this paper is organized as follows. In Section II, we describe the resource allocation model we consider, including cost structure and monetary transfers between users and nodes. In Section III, we precisely define the notion of stability and characterize the possible allocation results. In

Section IV, we present the best-response dynamics and the related convergence results. Some representative simulation results are reported in Section V. Section VI summarizes related work. Section VII concludes this paper and outlines future research directions.

II. Resource Allocation in Networks

A resource allocation game occurs in a *network*, which is given by a directed graph $G = (V, E)$, with network node set V and directed edge set E. There are two types of *agents* in our game: *users* and network *nodes*. The set of users that intend to route traffic is $\mathcal{A} = \{1, \ldots, k\}$. A user i wants to send a positive amount r_i of traffic from source node $s_i \in V$ to destination node $t_i \in V$. We use $\mathcal{P}_i$ to denote a set of paths from s_i to t_i, and user i intends to use one of those paths to route the traffic. We consider only networks in which $\mathcal{P}_i \neq \emptyset$ for all i, and define $\mathcal{P} = \cup_{i=1}^{k} \mathcal{P}_i$. We allow graph G to contain parallel edges, and a node can participate in multiple source–sink pairs. Each path $P \in \mathcal{P}_i$ consists of a set of ordered nodes starting from its source node s_i to its destination node t_i. For example, assume that user i wants to route 1 unit of traffic from node nd_1 to node nd_2. A path could be $nd_1 \rightarrow nd_3 \rightarrow nd_4 \rightarrow nd_2$. That is, if user i decides to use this path to route traffic, it needs to request node nd_3 to accept traffic from node nd_1, node nd_4 to accept traffic from node nd_3, and node nd_2 to accept traffic from node nd_4. Therefore, user i needs to make agreements with nodes nd_3, nd_4, and nd_2. We can see that user i does not need to make any agreement with its source node nd_1. For ease of analysis, we assume that a path does not include its source node. We also introduce the restriction that each node appears at most once in each path because a user only wants to route its traffic to a destination node, and, thus, visiting a node more than once cannot bring it any benefit.

The strategy of user i is a vector $\theta_i = \{b_{ij} : j \in V, s_i \neq j\}$. The entry $b_{ij} \geq 0$ represents a bid from user i to node j to request node j to route traffic for user i. The strategy of node j is a vector $\theta_j = \{o_{ij} : i \in \mathcal{A}, s_i \neq j\}$. The entry $o_{ij} > 0$ represents the minimum amount of payment that node j is willing to accept from user i to route traffic for user i. Given the strategies of all users and nodes, a contract between user i and node j is made iff $b_{ij} \geq o_{ij}$ and i only needs to pay $\psi(b_{ij}, o_{ij})$ to node j, where $\psi(b_{ij}, o_{ij})$ is a payment rule. A payment rule should satisfy the following: 1) $o_{ij} \leq \psi(b_{ij}, o_{ij}) \leq b_{ij}$, and 2) $\psi(b_{ij}, o_{ij})$ is nondecreasing with b_{ij} and o_{ij}. In this paper, payment rule $\psi(b_{ij}, o_{ij}) = o_{ij}$ is used when $b_{ij} \geq o_{ij}$, which implies that user i pays node j's demand, and node j only gets its demand.[1]

A path $P \in \mathcal{P}_i$ is *active* iff all nodes on the path have agreed to route traffic for user i. User i can send its traffic to its destination only if one of its paths is active. It is possible that more than one path of i is active. A user pays more if it makes agreements for more than one active path. However, its gain of delivering its traffic will not increase with more active paths. Here, we implicitly assume that each user has only one active path. Note that we consider an unsplittable flow problem in which user i is not permitted to fractionally route its r_i units of traffic over paths $\mathcal{P}_i$.

Given the strategy profile $\theta = \{\theta_i : i \in \mathcal{A} \cup V\}$ of all users and nodes, the set of active paths is $\mathcal{P}^a(\theta) = \cup_{i \in \mathcal{A}} \mathcal{P}_i^a(\theta)$, where $\mathcal{P}_i^a(\theta)$ is the set of active paths for user i. The amount of flow passing through or terminating at node j is $f_j(\theta) = \sum_{P \in \mathcal{P}^a(\theta), j \in P} f(P)$. $f(P) = r_i$ if P is a path for the s_i–t_i pair, i.e., $P \in \mathcal{P}_i$. Each node j has a cost function $c_j : \mathcal{R}^+ \rightarrow \mathcal{R}^+$. We always assume that cost functions are nonnegative, continuous, and nondecreasing. All of these assumptions are reasonable in applications where cost represents a quantity that only increases with the workload. In particular, we do not impose explicit node capacities. However, capacity constraints can be put in nodes' cost functions. For example, a node's cost goes up to infinity when the amount of traffic is higher than its capacity.

Next, we formalize the notion of utility and social welfare in the resource allocation game. User i's utility is defined as

$$u_i(\theta) = \begin{cases} -\sum_{j \in V, b_{ij} \geq o_{ij}} o_{ij} & \text{if } \mathcal{P}_i^a(\theta) = \emptyset \\ v_i(r_i) - \sum_{j \in V, b_{ij} \geq o_{ij}} o_{ij} & \text{otherwise} \end{cases}$$

where $v_i(r_i)$ is user i's gain when its traffic can be routed, and $\sum_{j \in V, b_{ij} \geq o_{ij}} o_{ij}$ is user i's cost for its contracts. User i's utility can be rewritten as $u_i(\theta) = v_i(r_i) \times \min(1, |\mathcal{P}_i^a(\theta)|) - \sum_{j \in V, b_{ij} \geq o_{ij}} o_{ij}$, where $|\mathcal{P}_i^a(\theta)|$ is the number of user i's active paths in $\mathcal{P}_i^a(\theta)$.

The utility of each node j is defined as

$$u_j(\theta) = \sum_{i \in \mathcal{A}, b_{ij} \geq o_{ij}} o_{ij} - c_j\left(f_j(\theta)\right).$$

Given all agents' strategies, the social welfare is defined as

$$\begin{aligned} sw(\theta) &= \sum_{i \in \mathcal{A}} u_i(\theta) + \sum_{j \in V} u_j(\theta) \\ &= \sum_{i \in \mathcal{A}} v_i(r_i) \times \min\left(1, |\mathcal{P}_i^a(\theta)|\right) - \sum_{i \in \mathcal{A}} \sum_{j \in V, b_{ij} \geq o_{ij}} o_{ij} \\ &\quad + \sum_{j \in V} \sum_{i \in \mathcal{A}, b_{ij} \geq o_{ij}} o_{ij} - \sum_{j \in V} c_j\left(f_j(\theta)\right) \\ &= \sum_{i \in \mathcal{A}} v_i(r_i) \times \min\left(1, |\mathcal{P}_i^a(\theta)|\right) - \sum_{j \in V} c_j\left(f_j(\theta)\right). \end{aligned}$$

Therefore, while keeping the set of active paths, changing agents' payments will not change the social welfare.

Each user has a set of paths that it can choose to route its traffic. Before using a path to route traffic, a user needs to make agreements with all the nodes on that path. The maximum amount of money user i can pay for making agreements for path $P \in \mathcal{P}_i$ is $v_i(r_i)$ (otherwise, the user gets negative utility by using this path). The minimum amount of money user i needs to pay for a node on the path P is the node's cost incurred by routing the traffic for user i. A node's cost in routing a traffic depends on its current traffic, which is determined by the strategies of other users.

[1]Network resource allocation games with other payment rules [e.g., $\psi(b_{ij}, o_{ij}) = b_{ij}$, $\psi(b_{ij}, o_{ij}) = (b_{ij} + o_{ij})/2$] can be analyzed analogously.

Definition 1 (Potentially Active Path): A path $P \in \mathcal{P}_i$ is potentially active if and only if there exists a strategy profile θ_{-i} of all agents without user i such that the total marginal cost $\sum_{j\in P}(c_j(f_j(\theta_{-i})+r_i)-c_j(f_j(\theta_{-i})))$ of all the nodes on path P while routing traffic for user i is no higher than user i's gain $v_i(r_i)$. Formally, a path $P \in \mathcal{P}_i$ is potentially active iff $\exists\theta_{-i}$, $v_i(r_i) \geq \sum_{j\in P}(c_j(f_j(\theta_{-i})+r_i)-c_j(f_j(\theta_{-i})))$.

If a path is not potentially active, its user does not need to consider it as an option for routing traffic. Without loss of generality, in this paper, we assume that all paths $\mathcal{P}$ are potentially active. While it is complex to check whether a path is potentially active, we introduce the notion of *strongly* potentially active path, which is easier to check. A path of a user is strongly potentially active if it is a potentially active path when there is only one user.

Definition 2 (Strongly Potentially Active Path): A path $P \in \mathcal{P}_i$ is strongly potentially active if and only if when there are no other users, the total cost $\sum_{j\in P} c_j(r_i)$ of all the nodes on path P while routing traffic for user i is no higher than user i's gain $v_i(r_i)$. Formally, a path $P \in \mathcal{P}_i$ is strongly potentially active iff $v_i(r_i) \geq \sum_{j\in P} c_j(r_i)$.

Proposition 3: If cost functions are concave and a path $P \in \mathcal{P}_i$ is strongly potentially active, it is also potentially active.

Proof: Given that cost functions are concave, for any θ_{-i} and any $j \in P$, it follows that $c_j(f_j(\theta_{-i})+r_i)-c_j(f_j(\theta_{-i})) \leq c_j(r_i)$. Thus, $v_i(r_i) \geq \sum_{j\in P} c_j(r_i)$ implies $v_i(r_i) \geq \sum_{j\in P}(c_j(f_j(\theta_{-i})+r_i)-c_j(f_j(\theta_{-i})))$. ■

Proposition 4: If cost functions are convex, a potentially active path $P \in \mathcal{P}_i$ is also strongly potentially active.

Proof: A potentially active path $P \in \mathcal{P}_i$ implies that $\exists\theta_{-i},\ v_i(r_i) \geq \sum_{j\in P}(c_j(f_j(\theta_{-i})+r_i)-c_j(f_j(\theta_{-i})))$. Since cost functions are convex, for any θ_{-i} and any $j \in P$, it follows that $c_j(f_j(\theta_{-i})+r_i)-c_j(f_j(\theta_{-i})) \geq c_j(r_i)$. Thus, $v_i(r_i) \geq \sum_{j\in P} c_j(r_i)$, i.e., P is strongly potentially active. ■

We mentioned before that the social welfare does not depend on agents' payments if their active paths remain the same. In other words, the social welfare only depends on users' active paths achieved by contracting. To optimize the social welfare, we need to find the set of paths that maximize the difference between users' gains by routing traffic and nodes' cost incurred by routing traffic. The social welfare maximization problem can be formulated as follows:

$$\mathcal{P}^a_{\text{opt}} = \max_{\mathcal{P}^a\subseteq\mathcal{P}}\left(\sum_{i\in\mathcal{A}} v_i(r_i)\times\min\left(1,|\mathcal{P}^a\cap\mathcal{P}_i|\right)-\sum_{j\in V} c_j\left(f_j(\mathcal{P}^a)\right)\right)$$

where $\mathcal{P}^a$ is a subset of $\mathcal{P}$, and $f_j(\mathcal{P}^a)=\sum_{P\in\mathcal{P}^a} f(P)$ is the amount of traffic passing or terminating at node j given the set of active paths $\mathcal{P}^a$. It is easy to see that $|\mathcal{P}^a_{\text{opt}}\cap\mathcal{P}_i| \leq 1$, i.e., in the optimal solution, each user has at most one active path.

Proposition 5: If cost functions are concave and a path $P \in \mathcal{P}_i$ is strongly potentially active, then $|\mathcal{P}^a_{\text{opt}}\cap\mathcal{P}_i| = 1$.

Proof: Assume that $\mathcal{P}^a_{\text{opt}}$ is an optimal solution and $|\mathcal{P}^a_{\text{opt}}\cap\mathcal{P}_i| = 0$. After adding path P to $\mathcal{P}^a_{\text{opt}}$, the utility increase of user i is $v_i(r_i)$, and the cost increase to all the nodes is $\sum_{j\in P}(c_j(f_j(\mathcal{P}^a_{\text{opt}})+r_i)-c_j(f_j(\mathcal{P}^a_{\text{opt}})))$. The change of social welfare is $\sigma_{sw}=v_i(r_i)-\sum_{j\in P}(c_j(f_j(\mathcal{P}^a_{\text{opt}})+r_i)-c_j(f_j(\mathcal{P}^a_{\text{opt}})))$. As each node's cost function is concave, it follows that $\sum_{j\in P}(c_j(f_j(\mathcal{P}^a_{\text{opt}})+r_i)-c_j(f_j(\mathcal{P}^a_{\text{opt}}))) \leq \sum_{j\in P} c_j(r_i)$. Thus, $\sigma_{sw} \geq v_i(r_i)-\sum_{j\in P} c_j(r_i) \geq 0$. Therefore, social welfare will increase by adding path P, which is contrary to the fact that $\mathcal{P}^a_{\text{opt}}$ is an optimal solution. ■

Theorem 6: The problem of computing the socially optimal allocation is $\mathcal{NP}$-complete and is impossible to approximate.

Proof: Membership of NP is easy. We generate a decision problem for the optimization problem first: Is there an allocation $\mathcal{P}^a_{\text{opt}}$ that can generate a social welfare at least α? Given an allocation $\mathcal{P}^a_{\text{opt}}$, compute the cost of each node and then sum up the cost of all nodes. Finally, compare the users' gains and nodes' cost. It is easy to see that the complexity of the calculation is polynomial in the number of nodes and the number of users.

For hardness, we must show that the problem is no easier than all other NP-complete problems. To do this, it suffices to show that any instance I of some known $\mathcal{NP}$-complete problems can be transformed into an instance of $\tau(I)$ of the optimization problem such that the transformation can be done in polynomial time, and the transformed problem $\tau(I)$ has a solution if and only if the original problem I has a solution. Here, we define a reduction from the arc-disjoint path problem, which has been shown to be strongly $\mathcal{NP}$-complete [11]. The disjoint path problem is specified by a graph G and a set of source–destination pairs $(s_1,t_1),(s_2,t_2),\ldots,(s_k,t_k)$. The goal is to find k disjoint paths $P_1,\ldots,P_k$ such that path P_i connects vertex s_i to vertex t_i. Given an instance of arc-disjoint path problem, we transform it into an instance of our allocation problem as follows. First, we create k users in which user i wants to send traffic from node s_i to node t_i. The amount of traffic of each user is one unit, i.e., $r_i = 1$ for all i. $v_i(r_i) = 1$. User i's path set $\mathcal{P}_i$ contains all paths from node s_i to node t_i given the directed graph G. The cost of each node j is

$$c_j(x)=\begin{cases}0 & \text{if } x <= 1\\ x\times\omega & \text{otherwise}\end{cases}$$

where $\omega > |E|$ is a large number.

If the initial instance has k arc-disjoint paths, the social welfare should be k if each user takes the corresponding path, as the traffic passing or terminating at each node should be at most 1. To see the other direction, if there exists an allocation with the social welfare of at least k, then there is a solution to the initial arc-disjoint path problem. Therefore, the allocation optimization problem is strongly $\mathcal{NP}$-hard.

In addition, the allocation optimization problem is inapproximable as we can set an infinitely large ω, which will make the distance between utility k and another utility arbitrarily large when two active paths are going through the same node. This implies that the allocation optimization problem cannot be approximated within a finite factor. ■

III. Characterization of Equilibrium

This section introduces the equilibrium concepts. The concept of NE is, in some sense, the analog of centralized optimal design in the context of multiple distributed selfish agents. However, multiplicity of Nash equilibria motivates an examination of a stronger equilibrium concept. Here, we also use pairwise NE [13], which is a refinement of NE in which a pair of agents is allowed to simultaneously change their strategies. The concept of pairwise NE fits our contracting framework in which a pair of agents is able to coordinate their actions through direct negotiation. Note that our equilibrium analysis focuses on agents' pure strategies.

A strategy profile θ forms an *NE* of a resource allocation game if for each agent i and each alternative strategy θ'_i

$$u_i(\theta) \geq u_i\left(\theta'_i, \theta_{-i}\right).$$

We say that an allocation (a set of active paths) $\mathcal{P}^a$ is *supported* via a given game relative to a profile of utility functions if there exists an equilibrium θ of the game such that $\mathcal{P}^a = \mathcal{P}^a(\theta)$. It is easy to see that in any NE, a user's proposal b_{ij} should be no larger than node j's asking price o_{ij} as, if so, the node can get higher utility by increasing o_{ij} to b_{ij}. In addition, the existence of NE is guaranteed.

Proposition 7: Nash equilibria always exist in all resource allocation games.

Proof: If any pair of user i and node j takes the strategy $b_{ij} = 0$ and $o_{ij} = \omega$, where ω is a large number, an empty allocation without any contract is formed. It is obvious that the empty allocation is always supported by an NE. ■

Proposition 8: The socially optimal allocation is supported by a strategy profile of NE.

Proof: Given the socially optimal allocation $\mathcal{P}^a_{\text{opt}}$, we generate agents' strategy profile as follows. If user i has no agreement with node j (i.e., user i's active path does not include node j), user i's payment to node j is 0, and j's demand for any user i is ω, which is a large number. Other b_{ij} and o_{ij} pairs need to satisfy the following conditions: 1) $b_{ij} = o_{ij}$; 2) if user i has an active path, its total payment is no higher than $v_i(r_i)$; and 3) if the amount of traffic passing or terminating at node j is positive, node j's received payment from all users is equal to node j's cost of routing traffic. It is easy to see that the above strategy profile exists and is an NE. Each node gets a utility of 0, and it cannot get a higher utility by deviating from the current strategy. Similarly, each user's utility will not increase by choosing a different strategy. ■

The *price of anarchy* (*price of stability*, respectively) is the worst (best, respectively) possible ratio of the social welfare found by some independent selfish behavior and the optimal social welfare possible by a centralized social welfare maximizing solution [25]. The price of anarchy and the price of stability are well-studied concepts in algorithmic game theory for problems such as load balancing, routing, and network design [25].

Proposition 9: The price of anarchy of the network resource allocation game could be arbitrarily large, and the price of stability of the network resource allocation game is 1.

Proof: The result is trivially followed by the fact that both the empty allocation and the socially optimal allocation are supported by a strategy profile of NE. ■

The empty NE allocation is formed because each user expects some others to make changes. It is easy to see that the concept of Nash stability is too weak as a concept for modeling contract formation when contracts are bilateral. Here, we use a stronger concept pairwise NE to accommodate bilateral interactions. This refinement allows any two agents who have not yet reached an agreement to change their bids to make an agreement.

Definition 10 (Pairwise NE): A strategy profile θ is a pairwise equilibrium of one of the above games if it is a NE of the game, and there does not exist any $i \in \mathcal{A}$ and $j \in V$ such that

$$u_i\left(\theta_{-ij}, b'_{ij}, o'_{ij}\right) > u_i(\theta), u_j\left(\theta_{-ij}, b'_{ij}, o'_{ij}\right) \geq u_j(\theta) \text{ or}$$
$$u_i\left(\theta_{-ij}, b'_{ij}, o'_{ij}\right) \geq u_i(\theta), u_j\left(\theta_{-ij}, b'_{ij}, o'_{ij}\right) > u_j(\theta)$$

where θ_{-ij} indicates the strategy profile found simply by deleting b_{ij} and o_{ij}.

Proposition 11: An NE is also a pairwise NE in the following cases: 1) each user has an active path; or 2) for any user i who has no active path, there is no path $P \in \mathcal{P}_i$ such that $|P| = 1$, where $|P|$ is the number of nodes in path P.

Proof:

Case 1) If user i and node j have an agreement, one agent's utility increase leads to the other agent's utility decrease when they simultaneously change the values b_{ij} and o_{ij}. If user i and node j have no agreement, the user's utility will decrease by making an agreement with node j.

Case 2) For user i who does not have an active path, it should have no agreement with any node in the NE. Therefore, its making a new agreement with a node cannot guarantee that it can route its traffic if there is no path $P \in \mathcal{P}_i$ such that $|P| = 1$. However, the user's cost will increase as it needs to pay to make an agreement. ■

Proposition 12: A pairwise NE always exists in each resource allocation game. The socially optimal allocation is also supported by a strategy profile of pairwise NE.

Proof: For each resource allocation game, there is at least one NE. Assume that a strategy profile θ is an NE but is not a pairwise NE. By definition, θ is not a pairwise NE only if there is a user node pair (i, j) such that one agent's utility will increase and the other agent's utility will not decrease if they change their strategies (b_{ij}, o_{ij}) to (b'_{ij}, o'_{ij}). It is easy to see that there is no contract between user i and node j (i.e., $b_{ij} < o_{ij}$). We can prove this by contradiction. If there is a contract between user i and node j (i.e., $b_{ij} = o_{ij} > 0$), we consider the following two situations: 1) if $b'_{ij} \geq o'_{ij}$, one agent's utility will decrease due to the deviation, which is contrary to the fact that both agents are willing to change to strategy pair (b'_{ij}, o'_{ij}); and 2) if $b'_{ij} < o'_{ij}$, one agent's utility will increase by decommitting from the contract, which is contrary to the assumption that θ is an NE in which no agent can benefit from unilaterally decommitting from a contract. Therefore, there is no contract between user i and node j under strategy profile θ, i.e., $b_{ij} < o_{ij}$.

As both agents are willing to change to use strategy (b'_{ij}, o'_{ij}), it follows that $b'_{ij} = o'_{ij} > 0$, which means that user i has no active path given strategy profile θ, as otherwise, its utility will decrease by making another contract. Therefore, user i makes no contract with all nodes in θ given that it has no active path given strategy profile θ. Thus, we can see that node j is user i's destination, as if not, user i still has no active path by making only one contract, and, thus, user i has no incentive to change its bid price to b'_{ij}.

Based on the above analysis, we can create a pairwise NE strategy profile θ' from an NE strategy profile θ as follows. If there is a user node pair (i, j) such that one agent's utility will increase and the other agent's utility will not decrease if they change strategies (b_{ij}, o_{ij}) to (b'_{ij}, o'_{ij}), where $b'_{ij} = o'_{ij}$, change their strategies to (b'_{ij}, o'_{ij}). After the strategy change, the new strategy profile should be an NE. This process stops until the evolved strategy is a pairwise NE. After user i's strategy is changed, user i will have an active path, which implies that user i will not have an incentive to concurrently change its strategy with any node. Therefore, this process will surely converge to a pairwise NE.

It follows from the above process that a socially optimal allocation is also supported by a strategy profile of a pairwise NE. In addition, a socially optimal NE allocation is also a pairwise NE. ■

In summary, NE and pairwise NE always exist in all resource allocation games. In some situations, there may be more than one NE and pairwise NE. The above analysis is based on the situation that all agents concurrently decide their strategies. However, given the multiplicity of (pairwise) Nash equilibria, agents have to choose from among these different equilibria, and the coordination problem faced by individuals is not entirely resolved. In addition, the calculation of (pairwise) Nash equilibria may have a high computation cost (even computational intractable), particularly when there are a large number of users and network nodes. This leads us to study the process by which individuals learn about the resource allocation game and revise their decisions on contracting over time. To study this issue, in Section IV, we introduce a best-response dynamics in which all agents negotiate with each other in discrete time, and agents take actions that are the myopic best response to other agents' current strategies. In addition, our analysis shows that both the empty allocation and the socially optimal allocation(s) are supported by strategy profiles of a (pairwise) NE. While studying best-response dynamics, we will investigate 1) how good the allocations evolved by the dynamics are and 2) what factors affect the convergence of the dynamics.

IV. Dynamics

This section studies the resource allocation under a natural form of myopic local best-response dynamics. The allocation of resources evolves in discrete steps, each consisting of three stages. During the first stage, a set of exogenously designated users can choose to decommit from contracts made before and send out new proposals to network nodes. Then, during the second stage, network nodes get back to users to indicate whether they are willing to make agreements. During the third stage, users can choose to decommit from contracts made before. At each stage, each agent takes myopic best-response actions to maximize its utility.

Let $\mathcal{C}_i^t$ be the set of agreements agent i has made at the beginning of round t. After a contract is decommitted, it will be removed from agents' contract set. After a contract is formed, it will be added into agents' contract set. Obviously, at time 0, we have $\mathcal{C}_i^0 = \emptyset$. For each agreement C made between node $V(C)$ and user $A(C)$, $\lambda(C)$ is the agreement price, and $T(C)$ is the time when the agreement was made. We consider a discrete-time dynamics that includes three stages at every round. At round t, q users $\mathcal{A}^t$ are chosen to propose to network nodes. Each user $i \in \mathcal{A}^t$ then will send a set of proposals to network nodes. Let the set of proposals received by node j be $\mathcal{W}_j^t$. At the second stage at round t, each node j will decide whether to accept each proposal $\langle i, b_{ij}\rangle$ from user i or not. If node j agrees to accept a proposal $\langle i, b_{ij}\rangle$, it sends a confirmation message to user i. Otherwise, it sends a rejection message to user i. Therefore, any agreement is based on the permission of both parties. At the third stage, each user $i \in \mathcal{A}^t$ can choose to decommit from some contracts in $\mathcal{C}_i^t$ if it is profitable to do so.Note that a user i is allowed to decommit at both stages 1 and 3, but its decommitment criteria are different: at stage 1, user i will decommit from an agreement with node j if it believes that it could get a cheaper agreement with node j; at stage 3, user i will choose one active path (if possible) based on its agreements and decommit from agreements not related to that active path. After the three stages at round t, the dynamics then continues to the next round $t + 1$ given the new allocation of resources and contracts made between users and network nodes. We assume that when a node is indifferent between accepting a proposal and rejecting a user's proposal (or keeping a contract or decommitting from a contract), the node will choose to accept the proposal (or keeping the contract). That is, each agent is benevolent while it is maximizing its utility.

An *activation rule* refers to a rule used for choosing the set of users $\mathcal{A}^t$ at each round t. A *uniform* activation rule chooses users $\mathcal{A}^t$ in the following way: The probability that a user i will be chosen to make proposals is q/k. That is, each user has an identical probability of being chosen.

The *state* of the dynamics at round t is defined by the set of contracts $\mathcal{C}^t = \cup_{i\in\mathcal{A}}\mathcal{C}_i^t = \cup_{j\in V}\mathcal{C}_j^t$. We say that $\mathcal{C}^t = \mathcal{C}^{t'}$ if and only if all the contracts in $\mathcal{C}^t$ are the same as the contracts in $\mathcal{C}^{t'}$, including the price of each contract. Given the set of contracts $\mathcal{C}_i^t$ of user i, i's active paths at time t are $\mathcal{P}_i^a(\mathcal{C}_i^t)$. The allocation at time t then is $\mathcal{P}^a(\mathcal{C}^t) = \cup_{i\in\mathcal{A}}\mathcal{P}_i^a(\mathcal{C}_i^t)$. Note that $\mathcal{C}^t = \mathcal{C}^{t'}$ implies that $\mathcal{P}^a(\mathcal{C}^t) = \mathcal{P}^a(\mathcal{C}^{t'})$, but not vice versa. The reason is that an allocation is only about the active paths for each user, but a contract also concerns the price of that contract.

Definition 13 (Convergence): Dynamics converge to an allocation if there exists t' such that, for $t > t'$, $\mathcal{P}^a(\mathcal{C}^t) = \mathcal{P}^a(\mathcal{C}^{t'})$. Dynamics converge to a state if there exists t' such that, for $t > t'$, $\mathcal{C}^t = \mathcal{C}^{t'}$. Dynamics converging to a state implies that dynamics converge to a stable allocation.

A. Myopic Best-Response Strategy

An agent taking myopic best-response strategy makes strategic decisions so that its utility is maximized at the end of each stage, and it will not consider other agents' concurrent actions at the same stage. At the beginning of each stage of each round, each agent has complete information about the status of any other agent. However, each agent does not know the concurrent actions of other agents within each stage.

At stage 1 of round t, each user $i \in \mathcal{A}^t$ can propose to network nodes. At the beginning of this stage, we let $\mathcal{C}_i^t = \mathcal{C}_i^{t-1}$ for any agent i, and there are two cases: $|\mathcal{P}_i^a(\mathcal{C}_i^t)| = 1$,[2] i.e., i has one active path that can be used to route traffic. The cost associated with the only active path P in $\mathcal{C}_i^t$ is $\sum_{C\in\mathcal{C}_i^t} \lambda(C)$. However, i can still send out requests to nodes to find an alternative active path that is less costly than $\sum_{C\in\mathcal{C}_i^t} \lambda(C)$. The other case is $|\mathcal{P}_i^a(\mathcal{C}_i^t)| = 0$, which implies that the user needs to make contracts to route traffic. At stage 1, each user acts as follows. First, user i checks the price of each contract and decommits from a contract C with node j if $c_j(f_j(\mathcal{C}^t)) - c_j(f_j(\mathcal{C}^t) - r_i) < \lambda(C)$, where $f_j(\mathcal{C}^t)$ is the amount of flow passing or terminating at node j given contract set $\mathcal{C}^t$. $c_j(f_j(\mathcal{C}^t)) - c_j(f_j(\mathcal{C}^t) - r_i)$ is the node j's marginal cost of routing traffic for user i if user i did not make an agreement with node j before. $c_j(f_j(\mathcal{C}^t)) - c_j(f_j(\mathcal{C}^t) - r_i) < \lambda(C)$ implies that it is possible that i can make an agreement with node j with a price lower than $\lambda(C)$. Let $c_{\min}^i$ be user i's cost for an existing active path, and it is given by

$$c_{\min}^i = \begin{cases} \sum_{C\in\mathcal{C}_i^t} \lambda(C) & \text{if } |\mathcal{P}_i^a(\mathcal{C}_i^t)| = 1 \\ v_i(r_i) & \text{otherwise} \end{cases}$$

where $|\mathcal{P}_i^a(\mathcal{C}_i^t)|$ is the number of active paths based on updated contract set $\mathcal{C}_i^t$. Then, for each path $P' \in \mathcal{P}_i$, user i first decides whether to request help from all the nodes in path P' by comparing the predicted cost $c_i(P') = \sum_{j\in P'} c(j)$ of the path P' and $c_{\min}^i$. The cost to get an agreement from node j is $c(j)$, and $c(j)$ is given by

$$c(j) = \begin{cases} \lambda(C) & \text{if } \exists C \in \mathcal{C}_i^t, V(C) = j \\ c_j\left(f_j(\mathcal{C}^t + r_i)\right) - c_j\left(f_j(\mathcal{C}^t)\right) & \text{otherwise} \end{cases}$$

where $c_j(f_j(\mathcal{C}^t + r_i)) - c_j(f_j(\mathcal{C}^t))$ is node j's marginal cost by routing traffic for user i, and we use this marginal cost as the price of making an agreement with node j. If user i has already made an agreement with a node j, i does not need to make an agreement with node j as i cannot benefit from decommitting from their contract and making a new contract $(c_j(f_j(\mathcal{C}^t)) - c_j(f_j(\mathcal{C}^t) - r_i) \geq \lambda(C))$. If $c_i(P') < c_{\min}^i$, user i sends out requests to get agreements for each node on path P'. Otherwise, user i will not try to make agreements for path P', and it needs to spend more than the cost of the existing active path. If user i wants to make a contract with node j in path $P' \in \mathcal{P}_i$ and it has not made an agreement with j, user i sends a proposal W to node j, and the price of the proposal is $\lambda(W) = c_j(f_j(\mathcal{C}^t + r_i)) - c_j(f_j(\mathcal{C}^t))$.

At stage 2, each node j decides if it decommits from some existing contracts in $\mathcal{C}_j^t$ and decides if it accepts or rejects some proposals from users. Let the proposals received by node j be $\mathcal{W}_j^t$. Each existing agreement can be treated as a proposal as a node can decommit from a contract and can make a contract by accepting a proposal. Formally, node j needs to choose a set of proposals $\mathcal{W}^*$ to maximize its utility, and the optimization problem can be formulated as

$$\mathcal{W}^* = \max_{\mathcal{W}\subseteq\mathcal{W}_j^t\cup\mathcal{C}_j^t} \left(\sum_{W\in\mathcal{W}} \lambda(W) - c_j\left(\sum_{W\in\mathcal{W}} r_{A(W)} \right) \right)$$

where $r_{A(W)}$ is the amount of traffic of the user that sends the proposal W.

Theorem 14: Each node's optimization problem at stage 2 of round t is $\mathcal{NP}$-complete.

Proof: Membership of NP is easy. We generate a decision problem for the optimization problem first: Is there a set of proposals $\mathcal{W}^*$ that can generate a utility at least α? Given proposal set $\mathcal{W}^*$, compute the cost for routing traffic and the sum of payment that the node can receive from the set of proposals. It is easy to see that the complexity of the calculation is polynomial in the number of proposals in $\mathcal{W}^*$. For hardness, we define a reduction from the well-known 0–1 Knapsack problem.

Formal definition of 0–1 Knapsack problem: There is a knapsack of capacity $c > 0$ and N items. Each item has value $v_i > 0$ and weight $w_i > 0$. Find the selection of items ($\delta_i = 1$ if selected, 0 if not) that fit, $\sum_{i=1}^N \delta_i w_i \leq c$, and the total value, $\sum_{i=1}^N \delta_i v_i$, is maximized.

To see how the reduction works, consider an instance of 0–1 Knapsack problem, and we transform it into our optimization problem as follows. Each item i corresponds to a proposal W_i. The value of each item corresponds to the payment of the corresponding proposal, i.e., $\lambda(W_i) = v_i$. The weight of each item corresponds to the amount of traffic associated with the corresponding proposal, i.e., $r_{A(W_i)} = w_i$. Node j's cost function is defined as

$$c_j(x) = \begin{cases} 0 & \text{if } x \leq c \\ \infty & \text{otherwise.} \end{cases}$$

It is easy to see that our optimization problem is equivalent to the original 0–1 Knapsack problem. The transformation is entirely automatic and is polynomial. Thus, each node's optimization problem of stage 2 is $\mathcal{NP}$-hard. ■

After computing the optimal set of proposals $\mathcal{W}^*$, node j will act as follows: 1) If a contract $C \in \mathcal{C}_j^t$ is not included in $\mathcal{W}^*$, node j will decommit from the contract and notify user $A(C)$ about it; 2) if a proposal $W \in \mathcal{W}_j^t$ is not included in $\mathcal{W}^*$, node j will notify user $A(W)$ about its rejection decision; and 3) if a proposal $W \in \mathcal{W}_j^t$ is included in $\mathcal{W}^*$, node j will add the proposal to its contract list and notify user $A(W)$ about its accept decision.

During the third stage at round t, each user i selects one active path (if there exists one) and decommits from other unnecessary contracts. More specifically, $\mathcal{P}_i^a(\mathcal{C}_i^t)$ is the set of active paths given by user i's contract set $\mathcal{C}_i^t$ after stage 2. If user i has no active path, it will decommit from all the existing contracts. If user i has some active paths, i.e., $|\mathcal{P}_i^a(\mathcal{C}_i^t)| > 0$, user i will choose the path $P \in \mathcal{P}_i^a(\mathcal{C}_i^t)$ with the lowest cost and maintain all the contracts related to this path, and, at the same time, decommits from all the agreements that are not related to path P.

[2]As a user can decommit from agreements at the end of each round, each user will not maintain more than one active path, as it will spend more by doing so.

B. Convergence Analysis

Here, we analyze the convergence property of the dynamics in which all agents take the myopic best-response strategy.

Proposition 15: In the following special cases, dynamics will converge to a pairwise stable NE, no matter how many users are allowed to make proposals at each round.

1) Users' path sets are not overlapping, i.e., for any path $P \in \mathcal{P}_i$ of user i and $P' \in \mathcal{P}_j$ of user j, $P \cap P' = \emptyset$.
2) All nodes' cost functions are linear, i.e., $c_j(x) = \beta_j \times x$.

Proof:

Case 1) This result is trivial as each node will receive at most one request. At the second stage of each round, all nodes will agree to users' request, and users will choose the cheapest path at stage 3. No user can find a cheaper path at later rounds.

Case 2) This result is also obvious. As cost functions are linear, a node's marginal cost of routing any more traffic does not depend on its current workload. Therefore, at the second stage, all nodes will agree to users' request, and users will choose the cheapest path at stage 3. No user can find a cheaper path at later rounds. ■

Proposition 16: If all nodes' cost functions are convex, i.e., $c'_j(x)$ is monotonically nondecreasing, and only one user is allowed to propose at each round, dynamics will converge to a pairwise NE.

Proof: After each user chooses its active path, no agent (including both users and network nodes) will decommit from the agreements for this path later. Take the user proposing at round 0 for example. At the end of round 0, the user will choose the path with the lowest marginal cost of all the nodes on that path. It is easy to see that the user cannot find another cheaper active path later given that 1) nodes' cost functions are convex, and 2) each node may have some other agreement(s). Similarly, the user proposing at round 1 will not change to another active path after round 1. Therefore, no agent will decommit from existing agreements later, and thus, the dynamics will surely converge. ■

Proposition 17: There are instances in which the dynamics converges to allocations that are not socially optimal, no matter which level of concurrence is chosen.

Proof: Such an example is shown in Fig. 1, where the cost function of each node is on the top of each node. Suppose that there are two users i and j both planning to route three units of traffic from node 1 to node 4. The socially optimal solution is that both users use the path passing node 2. However, the stable allocation generated by our dynamics can only be that both users choose the path passing node 3, no matter which concurrence level is chosen. The socially optimal allocation is not achievable unless each user considers the other user's concurrent behavior when $q = 2$. The reason for failing to reach the socially optimal stable allocation is that each user does not consider other users' concurrent actions at each round, which may lead to the final equilibrium being far worse than the socially optimal application. ■

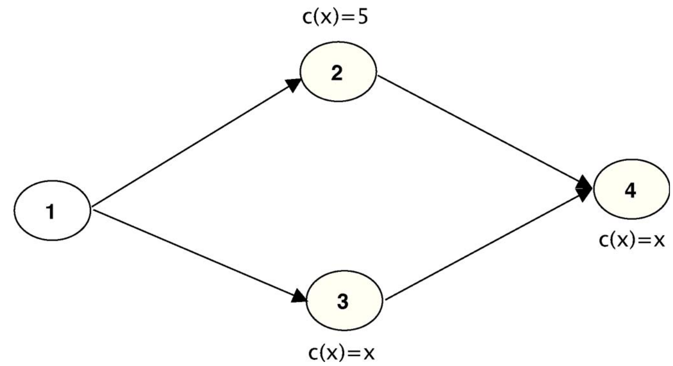

Fig. 1. Illustration example for inefficiency.

C. Concurrency and Decommitment Penalty

We consider several variations on our basic model of dynamics. The first variation we consider is concurrence, which is determined by the parameter q. If $q = 1$, only one user will propose to network nodes at each round. With the increase in the value of q, more users will concurrently propose at each round. The increase in concurrence will potentially increase the convergence rate of resource allocation, as more users can change their allocations at each round. However, on the other hand, too much concurrence may lead to more inconsistency on users' views about other agents. Thus, the increase in concurrence may slow down the convergence of allocation.

Another variation we consider here is the role of decommitment penalty. That is, when a user i wants to decommit from a contract C at round t, it needs to pay a penalty $\rho(C, t)$ to the node involved in the contract. Introducing penalty will add two features to our model: First, there is some probability that an individual exhibits "inertia," i.e., chooses the same strategy as in the previous period. This ensures that agents do not perpetually miscoordinate. Second, more commonly and realistically, there is a cost associated with making proposals, and penalty can be used to model such a cost. Costs can be incurred by the computation of various kinds of decisions. Costs can also be incurred by the pressure of a deadline by which a user has to stop negotiation. In addition, it is intuitive that one agent compensates the other contract party(ies) when it wants to be freed from a contract.

Once there is a penalty, an agent will change its strategy only if the gain from the deviation is higher than the penalty it needs to pay due to the deviation. In a resource allocation game with a lot of nodes and users, it may be very difficult to find a pure strategy (pairwise) NE. In the resource allocation dynamics, agents will dynamically choose contracting actions using myopic best-response strategy. A (pairwise) NE may not be actually reached given agents' concurrent myopic actions. Then, (pairwise) Nash equilibria may be too restrictive to characterize resource allocation dynamics. It is easy to see, in retrospect, that the use of approximate equilibrium is inherently needed. Penalty can be used to model approximate equilibrium. Assume that the penalty for the deviation is higher than ϵ. If each agent i's changing its strategy from θ_i to any θ'_i cannot improve the payoff by more than $u_i(\theta'_i, \theta_{-i}) - u_i(\theta) = \epsilon$, all the

TABLE I
VARIABLES

Variables	Values
Number of network nodes ($\|V\|$)	$[4, 15]$
Number of users (k)	$[2, 15]$
Amount of traffic (r_i)	$[5, 20]$
Concurrency rate (q/k)	$[0, 1]$
Number of paths per user	$[1, 8]$
η	$[0, 1]$

agents will choose not to deviate from their current strategies and the equilibrium is ϵ-approximate NE.

Given that it is difficult to get closed-form results, we now move to further simulations to get more intuitions about the behaviors of the best-response dynamics.

V. Numerical Results

We implemented a simulation testbed consisting of a manager, user agents, and network node agents to simulate the dynamics in the resource allocation game. The manager generates users and network nodes and randomly determines their parameters (e.g., the number of users, the number of network nodes, the level of concurrence, cost functions, the number of paths, and penalty rates). We tested the dynamics in different scenarios subject to those parameters. We characterize the pairwise stable Nash equilibria obtained with metrics such as dynamics' convergence rate, i.e., the social utility relative to the optimal social utility that is computed offline.

We performed a series of experiments in a variety of test environments, and the parameters are given in Table I. In the experiments, the number of network nodes in each graph is randomly selected from [4, 15]. The number of paths for each user is randomly selected from [1, 8]. The number of users is randomly selected from [2, 15], and each user can use any path to route its traffic. At each round, q users can propose concurrently. Concurrency rate q/k, which has a range of [0.1, 1], determines the ratio between the number of proposing users and the total number of users. A concurrence rate of 1 implies that, at each round, all the users make decisions concurrently. We use a linear decommitment penalty function $\rho(C, t) = \eta\lambda(C)$, where $\eta \in [0, 1]$. We ran each experiment for, at most, 500 rounds. If each agent's contracts does not change for over 50 rounds, which is long enough as compared with the number of users, we conclude that the dynamics of this experiment converge. The cost function of each node is in the form of $c_j(x) = \nu x^{\varrho}$, where $\nu \in [0.2, 5]$ and $\varrho \in [0.5, 1.5]$. Cost function $c_j(x)$ is 1) concave when $\varrho < 1$, 2) convex when $\varrho > 1$, and 3) linear when $\varrho = 1$. We found that varying the parameters of cost functions within a small range did not significantly affect the results. The amount of traffic associated with each user agent is randomly chosen from [5, 20].

After each experiment, we measure the social welfare of the final allocation. As we evaluate the dynamics in different environments, we measure the ratio of the utility of the evolved allocations to the utility of the optimal allocations. In each experiment, we use an exhaustive search to compute the socially optimal allocation. Another performance measure is the convergence rate, which is the ratio between the number of experiments in which dynamics converge and the total number of experiments.

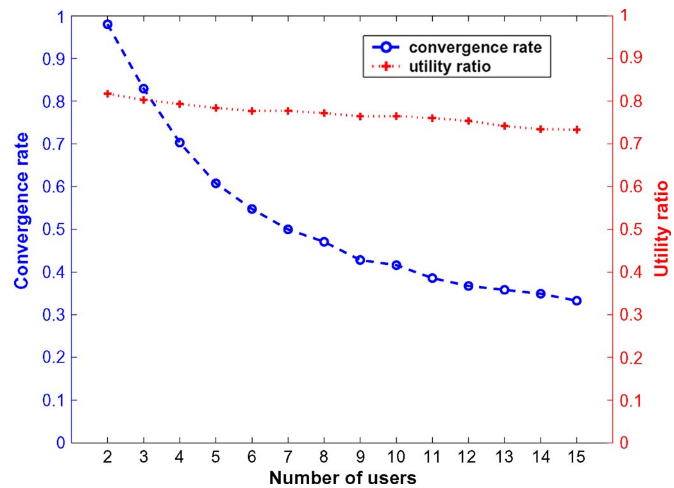

Fig. 2. Convergence rate and utility ratio as functions of the number of users.

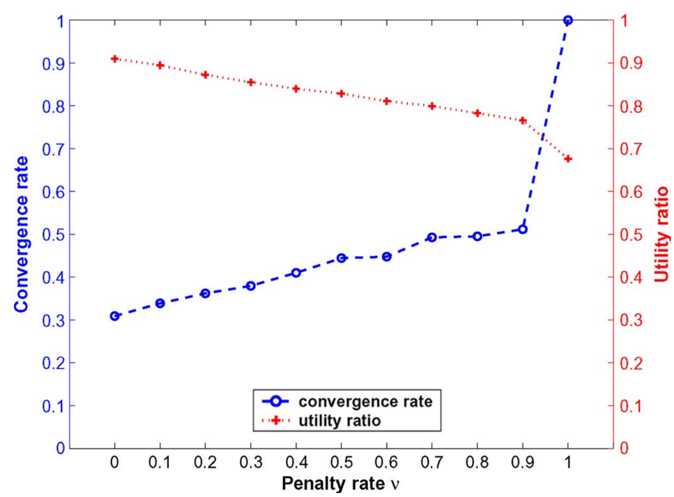

Fig. 3. Convergence rate and utility ratio as functions of the penalty rate.

Figs. 2–4 show some representative experimental results gathered from a series of experiments in different environment settings. Each data point represents the average value in 5000 runs where we randomly set the parameters in Table I.

Observation 1: The convergence rate and the utility ratio decrease with the increase in the number of users. Fig. 2 shows that when there are only two users, the convergence rate is almost 1. However, with the increase in the number of users, the convergence rate dramatically decreases, and the decreasing speed slows down with more users. This result is intuitive, as, with more users, competition will increase, and each user's marginal cost is harder to predict. One user's changing its strategy will result in more agents' strategy update. Fig. 2 also shows that with more users, the performance distance between the utility of the evolved final allocation and the optimal allocation increases, which corresponds to the intuition that it is more difficult to reach the socially optimal allocation when there are more users.

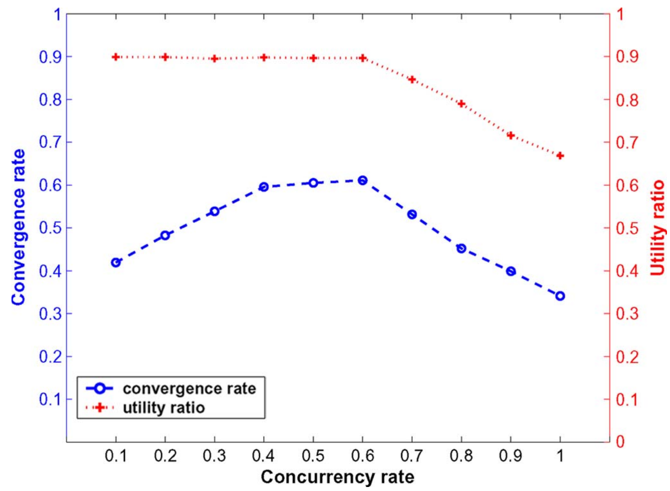

Fig. 4. Convergence rate and utility ratio as functions of the concurrence rate.

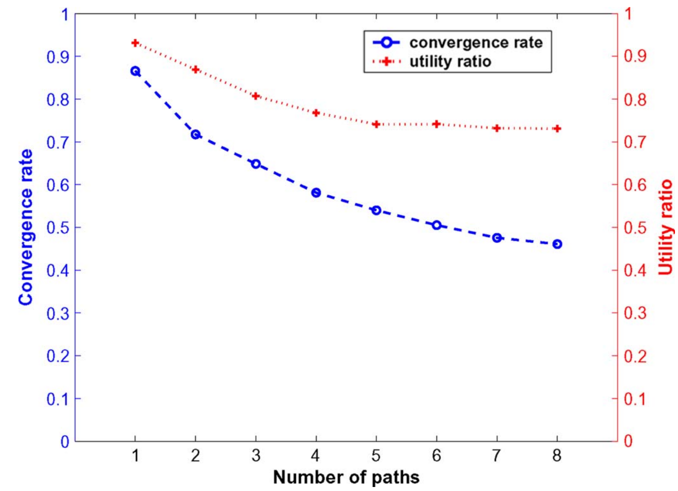

Fig. 5. Convergence rate and utility ratio as functions of the number of paths.

Observation 2: Experimental results in Fig. 3 suggest that the convergence rate increases with the increase in the penalty rate, which is intuitive as agents are more inclined to stick to existing contracts when the penalty is increased. We see that there is a big jump of the convergence rate when the penalty rate increases from 0.9 to 1. We also find that, with the increase in the penalty rate, the evolved final allocation becomes worse as compared with the socially optimal allocation, which corresponds to the intuition that, with higher penalty rates, users are making tradeoff between gaining better allocation and paying more penalties, and thus, agents' activity of "searching" for optimal allocations is decreased.

Observation 3: Experimental results in Fig. 4 show that the convergence rate increases with the increase in the concurrence when the concurrence is low, but the convergence rate decreases with the increase in the concurrence when the concurrence is high. When the concurrence is low, the increase in the concurrence may speed up the evolvement of dynamics, as more users can change their active paths at each round without too much interference. However, when many users are proposing concurrently, agents' strong dependence makes the dynamics difficult to converge.

From Fig. 4, we can also see that the utility of the evolved final allocation is getting worse when the concurrence is very high. That is because too many users' concurrent moves make it difficult for agents to move on the right path to the optimal allocation.

Observation 4: It can be observed from Fig. 5 that with the increase in the number of paths for each user, both the convergence rate and the utility ratio decrease. With more paths, each user has more alternatives to route its traffic. Accordingly, each user has more options while it decides to decommit from existing agreements. Thus, it is more difficult for the dynamics to converge. In addition, with more paths for each user, each network node can potentially be chosen by more users to route traffic. When a network node makes a new agreement or decommits from an agreement, it may result in the decommitment of other agreements.

Observation 5: Another observation we found from our experiments is that both the convergence rate and the utility ratio decrease when users' uncertainties increase. There are many sources of uncertainty in the resource allocation game. One such source is concurrence. In our myopic strategy, users do not reason about the actions of other concurrent proposing agents. As shown in Fig. 4, both the convergence rate and the utility ratio decrease when the concurrence is high.

In addition to concurrence, we also explored some other types of uncertainty in our experiments, such as if users do not know the set of contracts each node has made, users do not know other users' traffic load, and users have incomplete information about network nodes' cost functions. We observed that the convergence rate and the utility ratio decrease with the increase in uncertainty, which corresponds to the property in reality that agents often have better performance when agents have more knowledge of others.

VI. Related Work

Our research touches several lines of research. This section discusses some important related work like MARA, network formation, selfish routing, and automated negotiation.

Formally, a MARA problem can be represented as a triple $\langle \mathcal{A}, \mathcal{R}, \mathcal{U} \rangle$ representing sets of agents, resources, and utility functions by which individual agents associate values with resource subsets. MARA research is inspired by Sandholm's [27] work on sufficient and necessary contract types. Then, a number of variants of the MARA problem have been studied in the literature [4]. As in our approach, agents in MARA research (e.g., [4], [7]–[9], and [27]) often take myopic strategies, and they only accept deals from which they can perceive an *immediate* benefit [8]. However, our problem is different from the general MARA problem. In our problem, there are two types of agents: network nodes and users that have no resource. In MARA, an allocation of resources is partitioning of resources among all the agents [9]. In other words, each resource has at most one "owner." However, in our problem, each network node can route traffic for any number of users. That is, each "resource" can be used by multiple agents. In addition, network structures play an important role in the resource allocation results. Furthermore, MARA focuses on discrete resources, but in our problem, each node's cost function is continuous. The network resource allocation problem is much more complex than the general MARA problem.

Our work is also related to the work on network formation games by both the computer science and economics communities (see [12] and [31] for a survey). In our resource allocation problem, there are two types of agents: users that pay to have traffic routed and nodes in the network that accept payment and route traffic for users. In addition, an agent's utility depends on the routing result and the payments involved in agreements. However, in the network formation problem, the utility of each node is determined by the structure of the formed network and its efforts in forming the network, e.g., an agent incurs some costs while creating edges. In the model of Johari *et al.* [15], which also consider traffic routing through contracting, all the nodes form a network first, then each node takes the shortest path to send the traffic, and each node can use any node in the network. However, in our model, contracting is used to make agreements between users and network nodes, and the nodes will route traffic for the users only if they have made agreements.

This paper is also related to the literature on selfish routing [25], which is a classical mathematical model of how selfish users might route traffic through a congested network. The outcome of selfish routing is generally inefficient, and users always choose the path with the lowest cost (e.g., network delay). In contrast, in our paper, a user is always trying to minimize the cost of routing its traffic, while at the same time, each network node wants to maximize the difference between payments from the users and costs due to traffic routing. In addition, unlike congestion games in which costs incurred by a player depend only on its path and the amount of flow on the edges of its path [23], in our model, the cost is associated with each node, and the cost depends on the amount of traffic passing or terminating at that node.

This paper is also connected to the literature on automated negotiation that focuses on analyzing agents' strategic behaviors in various bargaining games characterized by the number of agents, the number of issues, and information setting (see [19] for a survey). In our problem, a user needs to negotiate for multiple resources (issues), and negotiation succeeds if the user gets agreements for one path. This is a simple form of multilinked negotiation, where the issues are interrelated in the sense that, from the perspective of the overall negotiation, negotiation issues are dependent, as an agent's utility from the overall negotiation depends on obtaining overall agreements. Almost all the work on multi-issue negotiation focuses on bilateral negotiation, and a variety of learning and searching methods are used, e.g., case-based reasoning [30], similarity criteria-based search [10], and decentralized search [17], [18]. The only exception is [2], which studies a multiresource negotiation in which resources belong to different agents. Our problem is more than a multiresource negotiation as there are multiple paths. One user just needs to get agreements for one path. Then, the user needs to decide which path to choose, and the optimal decision depends on many factors. To our best knowledge, there is no work investigating such a complex negotiation problem, and we leave it for future research. While automated negotiation work (e.g., [2], [10], [17], and [18]) is typically concerned with negotiation at the local level (what are agents' equilibrium strategies), in this paper, we address negotiation at a global level by analyzing how the actions taken by agents locally affect system-level objectives like social welfare or stability.

VII. Conclusion

The contributions of this paper include the following. 1) We have provided a framework for a multiuser resource allocation problem. We have shown that finding optimal resource allocation is $\mathcal{NP}$-complete and is inapproximable. We have taken both NE and pairwise NE as the solution concepts to characterize the equilibrium allocations. We have found that the social optimal allocations are supported by NE and pairwise NE. 2) A myopic best-response strategy for the resource allocation game has been proposed to characterize the resource allocation dynamics. We have also analyzed the convergence of the dynamics in some special cases. 3) Through experiments, we have shown that a) the dynamics' convergence rate i) decreases with the number of users, ii) increases with the penalty rate, iii) increases with the concurrence when the concurrence is low, iv) decreases with the concurrence after a certain level of the concurrence has been reached, v) decreases with the number of paths, and vi) decreases with users' uncertainty; and b) the ratio of the social welfare of evolved allocations and that of optimal allocations decreases with the increase in i) the number of users, ii) the penalty rate, iii) the concurrence, iv) the number of paths, and v) uncertainty.

There are several natural directions suggested by our research. The most obvious one is to expand the strategy space considered by each agent in our dynamics. For example, it would be interesting to develop some look-ahead strategies and more complex interaction protocols (e.g., infinite horizon strategic bargaining [26]). Again, agents' strategies should be natural and close to reality when we are trying to characterize the resource allocation game. Coordination (or mediation) mechanisms may be useful in optimizing the contracts of a set of users if they work together. In addition, our experiments, thus far, have focused on scenarios ranging from low to moderate complexity, but we wish to investigate much larger problems where there are more users and network nodes. In such cases, some new algorithms (e.g., approximate algorithms) are needed to compute the optimal allocation. In addition, in our model, a user needs to negotiate with all the nodes on a path. Future work will also consider a more complex model in which a user only needs to negotiate with the first node on a path, and the first node negotiates with other nodes to route the traffic.

The well-defined optimization problem considered in this paper relates to many applications (e.g., [1] and [28]). Our model may be a useful analytical tool for shedding light on the complex and dynamic processes that create the cooperative relationship in reality. Our work opens up the opportunity to consider contract-based resource allocation over large computational networks (e.g., supply chain, web/grid service composition, workflow, and enterprise integration). This paper assumes that all agents (including both users and nodes) have complete information about others. In more practical scenarios, agents only have incomplete information. In those cases, designing efficient or optimal mechanisms is a challenging problem. Analyzing agents' strategic behaviors under those mechanisms is also very interesting. This paper assumes that agents are selfish; it would also be interesting to study network resource allocation procedures where agents are cooperative (or semi-cooperative), particularly when agents have bounded communication and computational resources. It is also promising to design resource allocation dynamics in which each agent adopts strategies based on ideas from evolutionary game theory.

Acknowledgment

The authors would like to thank J. Kurose for his insightful comments that have significantly improved the quality of this paper. The authors would also like to thank the editors and the anonymous reviewers for their suggestions for improving the accuracy and readability of this paper.

References

[1] B. An, F. Douglis, and F. Ye, "Heuristics for negotiation schedules in multi-plan optimization," in *Proc. 7th Int. Joint Conf. Auton. Agents Multi-Agent Syst.*, 2008, pp. 551–558.

[2] B. An, V. Lesser, and K. M. Sim, "Decommitment in multi-resource negotiation," in *Proc. Int. Joint Conf. Auton. Agents Multi-Agent Syst.*, 2008, pp. 1553–1556.

[3] M. Branson, F. Douglis, B. Fawcett, Z. Liu, A. Riabov, and F. Ye, "CLASP: Collaborating, autonomous data stream processing systems," in *Proc. ACM/IFIP/USENIX 8th Int. Middleware Conf.*, 2007, pp. 348–367.

[4] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa, "Issues in multiagent resource allocation," *Informatica*, vol. 31, no. 1, pp. 3–31, 2006.

[5] S. H. Clearwater, *Market-Based Control: A Paradigm for Distributed Resource Allocation*. Singapore: World Scientific, 1996.

[6] P. Cramton, Y. Shoham, and R. Steinberg, Eds., *Combinatorial Auctions*. Cambridge, MA: MIT Press, 2006.

[7] P. E. Dunne and Y. Chevaleyre, "The complexity of deciding reachability properties of distributed negotiation schemes," *Theor. Comput. Sci.*, vol. 396, no. 1–3, pp. 113–144, May 2008.

[8] P. E. Dunne, M. Wooldridge, and M. Laurence, "The complexity of contract negotiation," *Artif. Intell.*, vol. 164, no. 1/2, pp. 23–46, May 2005.

[9] U. Endriss, N. Maudet, F. Sadri, and F. Toni, "Negotiating socially optimal allocations of resources," *J. Artif. Intell. Res.*, vol. 25, no. 1, pp. 315–348, Jan. 2006.

[10] P. Faratin, C. Sierra, and N. R. Jennings, "Using similarity criteria to make issue trade-offs in automated negotiations," *Artif. Intell.*, vol. 142, no. 2, pp. 205–237, Dec. 2002.

[11] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San Francisco, CA: Freeman, 1979.

[12] M. Jackson, "A survey of models of network formation: Stability and efficiency," in *Group Formation in Economics: Networks, Clubs, and Coalitions*. Cambridge, U.K.: Cambridge Univ. Press, 2004, pp. 11–88.

[13] M. Jackson and A. Wolinsky, "A strategic model of social and economic networks," *J. Econ. Theory*, vol. 71, no. 1, pp. 44–74, Oct. 1996.

[14] N. Jain, L. Amini, H. Andrade, R. King, Y. Park, P. Selo, and C. Venkatramani, "Design, implementation, and evaluation of the linear road benchmark on the stream processing core," in *Proc. 25th ACM SIGMOD Int. Conf. Manag. Data*, 2006, pp. 431–442.

[15] R. Johari, S. Mannor, and J. Tsitsiklis, "A contract-based model for directed network formation," *Games Econ. Behav.*, vol. 56, no. 2, pp. 201–224, Aug. 2006.

[16] F. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecommun.*, vol. 8, no. 1, pp. 33–37, Jan. 1997.

[17] G. Lai, C. Li, and K. Sycara, "Efficient multi-attribute negotiation with incomplete information," *Group Decis. Negot.*, vol. 15, no. 5, pp. 511–528, Sep. 2006.

[18] G. Lai, K. Sycara, and C. Li, "A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions," *Multiagent Grid Syst.*, vol. 4, no. 1, pp. 45–65, Jan. 2008.

[19] A. R. Lomuscio, M. Wooldridge, and N. R. Jennings, "A classification scheme for negotiation in electronic commerce," *Group Decis. Negot.*, vol. 12, no. 1, pp. 31–56, Jan. 2003.

[20] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds., *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.

[21] D. C. Parkes, "Iterative combinatorial auctions: Achieving economic and computational efficiency," Ph.D. dissertation, Univ. Pennsylvania, Philadelphia, PA, 2001.

[22] A. Riabov and Z. Liu, "Planning for stream processing systems," in *Proc. 20th Nat. Conf. Artif. Intell.*, 2005, pp. 1205–1210.

[23] R. W. Rosenthal, "The network equilibrium problem in integers," *Networks*, vol. 3, no. 1, pp. 53–59, 1973.

[24] M. Rothkopf, A. Pekec, and R. Harstad, "Computationally manageable combinational auctions," *Manage. Sci.*, vol. 44, no. 8, pp. 1131–1147, Aug. 1998.

[25] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. Cambridge, MA: MIT Press, 2005.

[26] A. Rubinstein, "Perfect equilibrium in a bargaining model," *Econometrica*, vol. 50, no. 1, pp. 97–109, Jan. 1982.

[27] T. Sandholm, "Contract types for satisfying task allocation: I. Theoretical results," in *Proc. AAAI Spring Symp.*, 1998, pp. 68–75.

[28] T. W. Sandholm and V. Lesser, "Coalitions among computationally bounded agents," *Artif. Intell.—Special Issue on Economic Principles of Multi-Agent Systems*, vol. 94, no. 1, pp. 99–137, Jul. 1997.

[29] T. Sandholm and V. Lesser, "Leveled commitment contracts and strategic breach," *Games Econ. Behav.*, vol. 35, no. 1, pp. 212–270, Apr. 2001.

[30] K. Sycara, "Problem restructuring in negotiation," *Manage. Sci.*, vol. 37, no. 10, pp. 1248–1268, Oct. 1991.

[31] E. Tardos and T. Wexler, "Network formation games and the potential function method," in *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.

Bo An (S'09) received the B.Sc. and M.Sc. degrees in computer science from Chongqing University, Chongqing, China, in 2003 and 2006, respectively. He is currently working toward the Ph.D. degree in the Department of Computer Science, University of Massachusetts, Amherst.

His research interests include autonomous agent and multiagent systems, particularly automated negotiation, resource allocation, and electronic commerce.

Victor Lesser (F'09) received the Ph.D. degree in computer science from Stanford University, Stanford, CA, in 1973.

He is currently a Distinguished Professor of computer science with the University of Massachusetts, Amherst, and is the Director of the Multi-Agent Systems Lab. His research focus is on the control and organization of complex artificial intelligence systems. He is considered a leading researcher in the areas of multiagent systems/distributed AI, intelligent distributed sensor networks, real-time AI, and blackboard systems. He has also made contributions in the areas of speech and sound understanding, diagnostics, adaptive signal processing, intelligent user interfaces, and computer architecture.

Dr. Lesser is a Founding Fellow of the American Association of Artificial Intelligence. He was the recipient of the 2009 International Joint Conferences on Artificial Intelligence Award for Research Excellence, was named a cowinner of the 2008 International Foundation for Autonomous Agents and Multi-Agent Systems (IFAAMAS) Influential Paper Award, and was also a recipient of a Special Recognition Award for his foundational research in generalized coordination technologies from the Information Processing Technology Office at the Defense Advanced Research Projects Agency. In 2007, The IFAAMAS established the "Victor Lesser Distinguished Dissertation Award" to recognize his career-long contributions to the field. He was the General Chair of the first International Conference on Multi-Agent Systems in 1995 and the Founding President of the IFAAMAS.