Fair Bandwidth Sharing in Distributed Systems: A Game-Theoretic Approach

Sudipta Rakshit and Ratan K. Guha, *Member, IEEE*

Abstract—Fair sharing of bandwidth remains an unresolved issue for distributed systems. In this paper, the users of a distributed LAN are modeled as selfish users with independence to choose their individual strategies. With these selfish users, the contention-based distributed medium access scenario is modeled as a complete-information, noncooperative game, designated the "Access Game." A novel MAC strategy based on p-persistent CSMA is presented to achieve fairness in the "Access Game." It is proven that there are an infinite number of Nash Equilibria for the "Access Game," but they do not result in fairness. Therefore, it may be beneficial for the selfish users to adhere to a set of constraints that result in fairness in a noncooperative fashion. This leads to the formulation of a constrained "Access Game" with fairness represented as a set of algebraic constraints. It is proven that the solution of the constrained game, the Constrained Nash Equilibrium, is unique. Further, it is shown that, in addition to achieving fairness, this solution also optimizes the throughput. Finally, these results are extended to a more realistic incomplete-information scenario by approximating the incomplete-information scenario as a complete-information scenario through information gathering and dissemination.

Index Terms—Distributed systems, local area networks, fair bandwidth share, selfish users, game theory.

1 Introduction

Bandwidth is one of the primary resources in computer communication networks and Quality of Service (QoS) is influenced by how bandwidth is allocated (in centralized systems) or shared (in distributed systems). As access to bandwidth is dependent on the Medium Access Control (MAC) protocol being used, the research on bandwidth fairness has focused on devising efficient MAC protocols to achieve fairness. A considerable amount of work has been done in this regard. Here, we briefly summarize it.

1.1 Previous Work

There are two types of MAC protocols: centralized and distributed. In centralized networks, fairness is based on the concept of Generalized Processor Sharing (GPS). Informally, GPS guarantees a user resource allocation proportional to that user's relative weightage [38]. We will follow this simple, yet powerful definition of fairness for the present work also. GPS cannot be implemented in practice because it relies on bit-by-bit switching, whereas the communication entity of interest is a packet. In [29], a practical packet-based implementation of GPS is presented. This algorithm is usually known as the Weighted Fair Queuing (WFQ) algorithm. In WFQ, each arriving packet is given virtual "start" and "finish" times based on the actual arrival time of the packet and the length of the packet. The packet with the smallest "finish" time is selected for transmission. A similar technique is presented in [39]. There, the WFQ algorithm designated Packet GPS (PGPS) is combined with a Leaky Bucket Admission Control algorithm for a single server GPS and it is shown that it is possible for the network to fulfill a wide range of performance guarantees using these algorithms. In [40], an improved GPS approximation algorithm, called Worst-case Fair Weighted Fair Queuing (WF2Q), is proposed. Using WF2Q, only packets with a virtual "start" time that has been passed are considered for transmission. This is a more accurate approximation of GPS, but increases implementation complexity. In distributed systems, there are two main problems in achieving fairness: lack of information, i.e., users usually do not know about the number of other users, and lack of coordination, i.e., one user cannot determine if any other is also transmitting simultaneously. Of these two problems, lack of coordination is more fundamental in nature because, even if the users know about the number of other users present in the system, medium access and packet transmission cannot be coordinated. Therefore, the objective of bandwidth fair sharing in a distributed system is to resolve this contention in such a way that users get bandwidth proportional to their weightages. For distributed systems, most of the work concentrates on Carrier Sense Multiple Access (CSMA) MAC protocols. In [35], a Distributed Fair Scheduling (DFS) scheme based on a virtual clock mechanism has been proposed for a Wireless Local Area Network (WLAN). As in WFQ, the "start" and "finish" times of an arriving packet are computed and the packet with the smallest "finish" tag is transmitted. A distributed algorithm using the back-off interval mechanism of IEEE802.11 MAC [11], [12] is used to determine the packet with the smallest "finish" tag. However, in general, it is difficult to implement a "virtual clock" mechanism in distributed systems. Another approach in differentiated bandwidth sharing in distributed systems is the "priority-based" access schemes. One of the earliest works incorporating priority in CSMA can be found in [41]. There, Tobagi presented a prioritized CSMA or P-CSMA.[1] The idea behind this scheme is as follows: grant access right exclusively to the messages of the current

• *The authors are with the School of Computer Science, University of Central Florida, Orlando, FL 32816. E-mail: {sraks, guha}@cs.ucf.edu.*

Manuscript received 11 June 2004; revised 10 Feb. 2005; accepted 31 May 2005; published online 16 Sept. 2005.
For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number TC-0197-0604.

1. Different from *p*-CSMA.

0018-9340/05/$20.00 © 2005 IEEE Published by the IEEE Computer Society

highest priority class. In [33], priority-based access schemes using the Carrier Sense Multiple Access (CSMA) protocol are analyzed for 1-persistent and nonpersistent modes. p-persistence CSMA is not considered due to the difficulty in analysis. Specifically, three schemes are considered: 1) all packets transmitted in 1-persistent mode, 2) higher priority packets transmitted in 1-persistent mode and lower priority packets transmitted in nonpersistent mode, and 3) all packets transmitted in nonpersistent mode. Assuming Poisson packet arrival and general packet length distributions mean packet delays are computed using approximate techniques. Another priority-based scheme for CSMA is presented in [34]. Some other related work for distributed systems can be found in [31] and [32]. The problem with the "priority-based" access schemes is the lack of an explicit relation between and priority and fairness. The medium access strategy we proposed is philosophically similar to the priority-based approach, but completely different in modeling and analysis.

1.2 Motivation and Contribution

The drawbacks of both the "virtual clock" and "priority-based" schemes have been indicated. Here, we present a novel approach for contention resolution by modeling the contention for medium access as a Non-Cooperative Game, the "Access Game." The "Access Game" model is predicated on an explicit relation between priority and fairness. Therefore, the solution of the "Access Game" satisfies the fairness definition as enunciated by GPS. Game Theory has been extensively used in other areas of computer communication [15], [16], [17], [18], [19], [20], [21]. However, there has been only limited application of Game Theory for designing distributed MAC protocols. It was recently introduced in [13]. To the best of the authors' knowledge, the present work is the first attempt to formulate the fairness problem in a Game-Theoretic framework. There are two reasons why Game Theory is a suitable tool for analyzing distributed medium access. First, the contention-based nature of medium access presents a natural application domain for Non-Cooperative Game Theory. Second, it is possible to conceive of "selfish" users in the future choosing their individual access strategies to optimize their own selfish interests [13], [37], [30]. The "virtual time"-based or "static priority"-based approaches described above are not suited for such situations. The "Access Game" model provides a theoretical formulation for achieving fair bandwidth sharing in the presence of "selfish" users. In addition to resolving the fairness problem, we also investigate in detail the interaction between the optimal "selfish" user strategies and the overall system performance.

1.3 Organization

The rest of the paper is organized as follows: Section 2 briefly describes some Game-Theoretic concepts and provides the "Access Game" model for the contention-based medium access, Section 3 analyzes a hypothetical complete-information scenario and proves the uniqueness of transmission strategies achieving fairness, Section 4 provides techniques to approximate an incomplete information scenario to a complete information scenario and shows that fair bandwidth share is achieved in this approximate case also, and Section 5 concludes this paper.

2 Modeling the Access Game

Formally, a finite game G consists of a nonempty finite set I of players. A player, say i, has a set of possible strategies/actions A_i. In order to play the game, all the players choose an action from the respective strategy sets simultaneously. At the end of the game, there is an outcome. Clearly, the outcome space is given by $S = \times_i A_i$. Let $s \in S$ be a generic outcome of the game G. Associated with the outcome, s, is a payoff to each of the players. Let us designate by $u_i = u_i(s)$ the payoff function for the ith user. The payoff function of the game is an ordered tuple of payoffs to individual users and is given by $u(s) = (u_1(s) \ldots u_n(s))$. Initially, Game Theory concentrated on games of pure strategy, i.e., user i taking only one action with probability "1" to play the game. Nash introduced the concept of mixed strategy. The concept of "mixed strategy" is that, instead of deciding for a particular action with certainty, a user i randomizes its decision and chooses a particular action from A_i with a probability (may be zero). Consequently, the elements of the outcome set S also become probabilistic in nature. As the payoffs are associated with the outcome of the game, it follows that, in a mixed strategy game, there is a nonnegative probability attached to the value of the payoff a user receives by playing the game. This entails the formulation of utility function $\bar{u}_i$ for each player i. In a mixed strategy game G, utility function $\bar{u}_i$ is the expected payoff for player i. We consider two solution concepts of a Non-Cooperative Game: Nash Equilibrium [26] and Constrained Nash Equilibrium [28]. For detailed discussion on these topics, see [14].

2.1 Access Game

Examples of distributed MAC protocols are ALOHA [3], Carrier Sense Multiple Access/Collision Detection (CSMA/CD) [5], [6], [7], [8], [9], and Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) [10], [11], [12], whereas an example of centralized MAC protocol is HIPERLAN [1], [2]. For the present work, a p-CSMA type access strategy for medium access is considered. Before proceeding further, a brief description on the p-CSMA protocol is in order. Consider a MAC protocol using CSMA. When a user has a packet to send, it senses the medium and, if the medium is sensed busy, it can take several actions and, depending on these actions, CSMA protocols can be classified as follows [4]:

1. 1-persistent: Keep on sensing the medium and transmit the packet when the channel becomes idle.
2. Non-persistent: Do not sense the medium for some time (i.e., back-off).
3. p-persistent: Keep on sensing the medium as long as the medium is busy and when the medium becomes idle, transmit with probability p and wait with probability (1-p).

We provide the following diagram (Fig. 1) for a schematic representation of p-CSMA.

It can be observed from Fig. 1 that, at the end of a transmission period, there is a brief idle period. The idle period essentially signals the end of the previous

Fig. 1. Successive states of the system in *p-CSMA*.

transmission period. After this brief idle period, users contend with each other to access the medium. Contention is eventually resolved in one user's favor and the successful user transmits next.

The MAC strategy presented for the "Access Game" is similar to the *p*-CSMA. The significant difference is that the value of "p" is not constant in the proposed protocol. Users compute appropriate values of p, based on the state of the system. More specifically, users transmit with different probabilities depending on the number of users present in the system. These are the players of the "Access Game." At the beginning of each contention period, each player has two actions to choose from: "transmit" and "wait." User i transmits with probability p_i. The game has three outcomes: "success," "failure," and "waste." The outcomes and payoffs of the game are as follows:

1. If no user transmits, the game's outcome is "waste" and all the users receive a payoff of "0."
2. If exactly one user transmits (say user i), then the outcome of the game is "success." User i receives a payoff of "1" and all the other users receive a payoff of "0."
3. If more than one user transmits, collision occurs and the outcome of the game is "failure." Every user receives a payoff of "0."

Note that if a user decides to "wait," irrespective of the game's outcome, it receives a payoff of "0." This choice of payoff is justified by the selfish nature of the users. Let us give a simple example regarding the payoff: Consider two users A and B contending for transmission. If both of them transmit, there is a collision and nobody benefits; hence, a payoff of "0" is assigned to both the users. If A (or B) decides to wait, it may so happen that B (or A) transmits and benefits. However, no benefit is accrued to A (or B). Therefore, A (or B) receives a payoff of "0." The payoff of "1" for success is self-explanatory.

Following the above discussion, the utility function or the expected payoff function can be expressed as:

$$\bar{u}_i = \Pr(success) \times 1. \qquad (1)$$

From here on, we drop the "bar" in $\bar{u}_i$ and denote the payoff function of user i as simply u_i. The main results are now presented in the following section.

2.2 Social Welfare

We have modeled the medium access as a Non-Cooperative Game. In the next section, solutions to the "Access Game" are provided. Before presenting these solutions, we briefly discuss some desirable qualities of these solutions in terms of Social Welfare and Pareto Optimality. Social Welfare (SW) is a concept from economics dealing with the distribution of resources in the society. Interestingly, this concept can be applied in the present context of medium access also, with the resource in question being the bandwidth. We have previously argued that each user should receive their fair share of bandwidth. Therefore, SW for the "Access Game" is achieved if these fairness criteria are satisfied. Associated with the concept of social welfare is the concept of Pareto-Optimality (PO) of a resource allocation. Informally, PO is achieved if nobody can be made better off without making somebody else worse off. All PO solutions do not result in SW, but all the allocations resulting in SW are PO. Whether the PO solutions satisfy the condition of SW depends on how SW is defined in a particular context. As the solutions in the next section satisfy SW within the present context, the solutions are PO in nature as well.

3 Analysis of the Access Game

From the discussion in the previous section, it is clear that each user has only one decision variable: the probability to transmit. In this section, we present the main results regarding these decision variables. First, the assumptions for the analysis conducted in this section are specified.

3.1 Assumptions

A1. Each user has complete information about all the other users. Although this assumption is quite restrictive for a distributed system, it is made to provide a sharp analysis. This assumption is later relaxed.

A2. All the users have packets to transmit. This assumption is made for the sake of simplicity. We also relax this assumption in the next section.

A3. Packets are of equal length. This assumption is again made for simplicity.

A4. The number of users playing the game is n and this number does not change. As $n = 1$ presents a trivial case, $n > 1$ is assumed.

A5. The system is stable.

Assumptions 4 and 5 are related. In this section, we present the results for a fixed number of users. In the next section, we relax this assumption and allow for varying numbers of users. For our analysis to hold for the general case in the next section, the system needs to be stable. The stability assumption has been maintained throughout this paper.

3.2 Nash Equilibrium

We first prove that the Nash Equilibria for the "Access Game" is inefficient.

Theorem 1. *In the Nash Equilibrium for the Access Game, there is at least one user i such that $p_i^{NE} = 1$.*

Proof. Reproducing the utility function of user i from (1), we have

$$u_i = p_i \prod_{j \neq i}^{n} (1 - p_j). \qquad (2)$$

Of interest is the expression $\prod_{j \neq i}^{n} (1 - p_j)$. Let us call it δ_i:

$$\delta_i = \prod_{j \neq i}^{n} (1 - p_j). \qquad (3)$$