Efficient Contention Resolution Protocols for Selfish Agents

Amos Fiat* Yishay Mansour† Uri Nadav‡

September 29, 2006

"Alright people, listen up. The harder you push, the faster we will all get out of here."

Police Chief Wiggum to crowd in post office
at tax filing deadline
— The Simpsons

Abstract

We seek to understand behavior of selfish agents accessing a broadcast channel. In particular, we consider the natural agent utility where costs are proportional to delay. Access to the channel is modelled as a game in extensive form with simultaneous play.

Standard protocols such as Aloha are vulnerable to manipulation by selfish agents. We show that choosing appropriate transmission probabilities for Aloha to achieve equilibrium implies exponentially long delays. We give a very simple protocol for the agents that is in Nash equilibrium and is also very efficient — other than with exponentially negligible probability — all n agents will successfully transmit within cn time, for some small constant c.

1 Introduction

Ethernet buses and wireless communications are both examples of shared communication media. Transmission is successful on such channels only if exactly one user accesses the media. Should multiple users access the channel simultaneously, a collision is said to occur, and all attempted transmissions fail. Contention resolution protocols are designed to address the problem of collisions, and to ensure fair and efficient use of such channels.

One would like to have a distributed contention resolution protocol, where anonymous users know little, if anything, about others. The celebrated Aloha protocol is an excellent example of such a distributed contention resolution protocol. Since the introduction of the Aloha protocol, much research has been devoted in deriving improved contention resolution protocols, where the main emphasis has been the stability of the protocol at high loads. (See [9] for an excellent treatment of the topic.)

Assume n agents at time zero, each with one packet to transmit. Agents that transmit without collision on the channel are said to be *successful*. A successful agent departs and is no longer in contention for the channel. If each of k yet unsuccessful agents transmits with probability $1/k$, then the expected latency per agent is $\Theta(n)$, and with high probability no agent will be unsuccessful after $10n$ time slots. For symmetric protocols this is the *socially optimal protocol* in terms of minimizing the expected sum of latencies, expected maximum, etc.

In this work we study contention resolution in the context of selfish user behavior. In the problems we study, rational and selfish agents seek to minimize their own costs, and have no compulsion to avoid harming others. One simple example of an agent's cost function is the *latency cost*, where the cost of a packet is the time delay between packet creation and successful packet transmission. Unless stated explicitly otherwise, we consider latency costs hereinafter.

Rational selfish agents with latency costs will subvert the socially optimal protocol given above. Consider Alice who continuously transmits until successful. If the other agents follow protocol and transmit with probability $1/k$, then the expected latency for Alice drops from $\Theta(n)$ to $O(1)$.

One can view the problem of devising protocols for selfish agents as a problem in mechanism design. However, we stress that we only allow protocols that are

*School of Computer Science, Tel Aviv University. Supported in part by a grant from the German-Israel Foundation. Email: fiat@cs.tau.ac.il.

†School of Computer Science, Tel Aviv University. Email: mansour@cs.tau.ac.il. The research was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a grant no. 1079/04 from the Israel Science Foundation, by a grant from BSF and an IBM faculty award. This publication only reflects the authors' views.

‡School of Computer Science, Tel Aviv University. Supported in part by a grant from the German-Israel Foundation. Email: urinadav@cs.tau.ac.il.

self-enforceable and do not involve external payments or incentives. In fact, we view protocol design here as searching for "good" equilibria, with seemingly surprising results. Using the terminology of [5], what we prove here is that the price of anarchy for contention resolution games is infinite, whereas the price of stability[1] for contention resolution games is $O(1)$.

A priori, one might naturally suspect an impossibility result, that all agents will continuously transmit, and therefore that no success will ever occur. Many examples of such selfish behavior have been shown in the game theory literature, this includes the prisoners dilemma and the "tragedy of the common", where players need cooperation to profit from a common resource.

In our setting, if the number of agents is at least three, then continuous transmission by all agents is indeed in equilibrium. In this natural equilibrium no transmission is ever successful. We call such an equilibrium is called a blocking equilibrium. However, there exist other *non-blocking* equilibria, where — eventually — all agents succeed. To see why this is so, consider two agents, Alice and Bob, each with one packet to send, and both seeking to minimize latency cost (delay until packet transmission). If Alice chooses to be aggressive and broadcast endlessly until successful transmission, the best response for Bob is to allow Alice to transmit in the first time slot, following which Alice loses interest in competing with Bob, and Bob now has full access to the channel without further interference.

We generalize the Alice and Bob example above to multiple agents, Alice, Bob, Carol, ..., and again — assume one packet per agent. A strategic player, Bob, will somehow have to balance the following profit/loss outcomes as influenced by his actions:

1. Immediate success, if Bob choose to transmit and none other did so, or

2. Delayed gratification, if Bob refrained from transmission and some other other agent was successful. Bob has gained because there are now less agents in contention for the shared media, or

3. "Wasted" time slots: either collisions or no transmissions on the channel, neither Bob nor any other agent is successful.

It may be illuminating to contrast our game with the repeated Prisoners Dilemma. In the finite horizon repeated Prisoners Dilemma, defecting is always the right choice. In our contention resolution setting this is not true. In a repeated game, the next game is exactly the same as the current game irrespective of the outcome of the current game. In our setting of a simultaneous play exhaustive form game, the game to be played next depends on the outcome of the current game. Fortunately, this gives non-trivial and socially desirable equilibria for various utility functions, even given fixed predetermined horizons.

Furthermore, it seems that the "folk theorems" [7] about cooperation and punishment in repeated games are not directly applicable to our problem. In the folk theorems, a misbehaving player may gain momentarily, but will receive punishment soon thereafter. In out setting, a defecting player may succeed in attaining his ultimate goal via defection (for example, hogging the transmission channel until successful transmission). Unlike a repeated game, defectors who have concluded their affairs will not hang about to receive punishment. This said, our efficient protocols in equilibrium do include the introduction of fear from overhanging "global" disaster that induces what seems to be cooperation.

We deal with synchronous communications where transmissions are only possible in discrete time slots. We assume that every agent has one packet to transmit[2]. Additionally, we deal with symmetric protocols, where all agents follow the same set of rules[3]. Our primary interest is in latency costs as agent utilities, but we also study the effect of a deadline.

We study the effect of time dependency on protocols. Time-independent protocols may determine transmission probability using the current number of agents in contention (called pending agents), but may not use the current time slot index. A time-dependent protocol is not so restricted. For example, the socially optimal protocol (prob. $1/k$ for k pending agents) is a time independent protocol but is not in equilibrium.

We show that there is a unique time-independent non-blocking symmetric protocol in equilibrium, in which all agents broadcast with probability $p_k \in \Theta(1/\sqrt{k})$ (again — k is the number of pending agents). With such transmission probabilities, the expected duration until all n agents succeed is approximately

[1]With one major caveat, that this only holds with high probability.

[2]Although our results also hold if agents start with multiple packets.

[3]In applications such as mobile applications over broadcast channels, anonymity occurs naturally and may even be a requirement. In such settings, it makes little sense to consider non-symmetric protocols where the three agents Alice, Bob, and Carol, each play a different strategy depending on their own identity and the identities of others with which they are playing. If we were to allow known (and unique) identities then the contention resolution problem becomes somewhat uninteresting. One could use social rank to determine priorities based upon identities, and this result is in equilibrium (possibly considered unfair by those of lower social standing).

$\sqrt{n}e^{\Theta(\sqrt{n})}$ (which is dominated by the expected latency of the first successful transmission). We deal with time independent protocols for latency in Section 5.

We define a protocol to be *efficient* if the maximal packet latency is linear with high probability. The socially optimal protocol (which sends with probability $1/k$) is efficient but is not in equilibria. Broadcasting with probability $p_k \in \Theta(1/\sqrt{k})$ is in equilibria but is not efficient. Furthermore, our claims above imply that any protocol that is simultaneously efficient and in equilibria must be time-dependent. Thus, we seek an efficient time-dependent protocol in equilibrium for latency costs. This motivates our study of deadline cost functions and suggests the notion of virtual deadlines, which we can use to derive efficient protocols.

A deadline cost function would typically charge only those agents that have not been successful prior to the deadline. *E.g.*, a tax audit for those not filing by midnight. Perhaps surprisingly, one sees dramatic behavioral changes in equilibria as a function of the time left until the deadline. If the deadline is close by (say $2n$ time slots away), then the only equilibria for selfish agents is to transmit with high probability (and thus the probability that *any* agent will be successful is negligible). Given a deadline $10n$ time units away, then — with very high probability — all n agents will succeed prior to the deadline. Deadline cost functions can be used to model Quality of Service issues, *e.g.*, MPEG packet delivery past a deadline causes video breakup. Section 6.1 deals with deadline cost functions.

We seek equilibria[4] where "ill behaved" latency cost agents behave more like "polite" deadline cost agents, for an appropriately chosen deadline. We stress again that we are not introducing external payments or charges to introduce the deadline, and we are not changing the latency cost assumption about the agents. Our protocol is "self policing" and enforces a "virtual" deadline on the agents, of sufficiently great cost so that they transmit with low probability $O(1/n)$.

2 Related Work

Altman *et al.* [1, 2], study a game theoretic model of slotted Aloha. In their work a very realistic model is studied, where agents have incomplete information as to the number of agents pending. They also assume a stochastic arrival flow to each source. In [2] agents' objective is delay minimization and in [1] agents' objective is to increase their throughput. Agents' strategies are restricted though, to a single retransmission probability. They show the existence of an equilibrium and give a numerical analysis of the model that shows that the system is inefficient by increasing the delays unduly, even under light traffic.

MacKenzie and Wicker [6] study stability of slotted Aloha, with selfish agents in the multi-packet reception model. They assume that agents utility is a function of the number of attempted transmissions before success, (*e.g.*, costs reflect power lost per successful transmission). They show the existence of equilibrium strategies in this model. They also show that for specific parameters, there exists points of equilibrium that attain the maximum possible throughput of Aloha.

There has been very extensive work on routing in networks by selfish agents (see [10] for an extensive survey). Much of this work is a study of Waldrop equilibria for traffic, cast as as either a network routing problem or as a machine load balancing problem. One seeks to understand equilibria for multicommodity flow where edges used by the flow introduce "latency", this latency is a function of the magnitude of flow through the edge. In contrast we are interested in the completion time of the task (in our case, sending a packets over a shared media).

We also mention that previous work has been done on distributed mechanisms [3]. One example of distributed mechanism design is the BGP protocol of [4]. This protocol makes use of side payments to ensure truthfullness, and computes these side payments with a distributed algorithm. Unfortunately, the algorithm itself is not incentive compatible.

3 The Model and Definitions

We consider the following contention resolution problem. Consider a set of n agents, each of which has a *single* packet to transmit. Agents that have not yet successfully transmitted their packet are called *pending*, initially all n agents are pending.

Time consists of discrete time slots. Agents that are pending at time slot t can either *'Transmit'* or be *'Quiescent'*. If exactly one agent chooses to transmit at time slot t then this agent is *successful* and ceases to be pending. If multiple agents choose to transmit at time t then a *collision* occurs. In case of collision or if the channel is idle then the set of pending agents remains unchanged. The number of agents at time zero, n, is known to all agents, and the agents keep track of K_t — the number of pending agents at time t.

We study multiple agent access to a channel as a non-cooperative game in extensive form and simultaneous play. The latency T_i for agent $i \in \{1, \ldots, n\}$ is a random variable whose value is the time at which agent i is successful (or ∞), and whose distribution is determined by the (possibly) mixed strategies of the agents.

[4] We remark that our protocol is not only in equilibria but also subgame perfect.

The cost to agent i is a function of the latency, $\Gamma(T_i)$, and is thus also a random variable. Our primary interest is in the latency cost function for all agents, *i.e.*, $\Gamma(t) = t$. We also present results for deadline cost functions (*e.g.*, $\Gamma(t) = 0$ for $t < D$ and $\Gamma(t) = M$ for $t \geq D$).

Definition 3.1. *A strategy for agent Alice, $q = \langle q_{k,t} : 1 \leq k \leq n, 0 \leq t \rangle$, is interpreted as follows: if Alice is one of k pending agents at time t (i.e., $K_t = k$), then Alice transmits with probability $q_{k,t}$.*

A strategy for agent Alice is said to be time-independent *if the transmission probabilities, $q_{k,t}$, are independent of the time, i.e., $q_{k,t} = q_{k,t'}$, for all $0 \leq t, t'$. A time-independent strategy can thus be represented as a vector $q = \langle q_1, q_2, \ldots, q_n \rangle$, where q_k is the transmission probability given k pending agents, irrespective of the time*[5].

Definition 3.2. *A* protocol *$Q = \langle q^{(1)}, q^{(2)}, \ldots, q^{(n)} \rangle$ is a list of strategies, one per agent, where agent i, $1 \leq i \leq n$ has strategy $q^{(i)}$.*

Fix a protocol Q. We define the expected cost of the protocol for agent i, $C_i^Q = E[\Gamma(T_i)]$. The expectation is taken over the probability distribution defined by Q. Let $T_{i|k,t}$ denote the latency for agent i, conditioned on $K_t = k$, and on agent i being one of the k pending agents. Let $C^Q_{i|k,t} = E[\Gamma(T_{i|k,t})]$, and define the expected future cost $F^Q_{i|k,t} = C^Q_{i|k,t} - \Gamma(t) = E[\Gamma(T_{i|k,t})] - \Gamma(t)$. (When clear from the context we drop the superscript Q.)

Definition 3.3. *Let $Q = \langle q^{(1)}, \ldots, q^{(n)} \rangle$ be a protocol. Let (s, Q^{-i}) denote the protocol where agents $j \neq i$ use strategies $q^{(j)}$ and agent i uses strategy s. We say that strategy s is a* best response *of agent i to Q^{-i}, if the expected cost to i with s, given that other agents $j \neq i$ use $q^{(j)}$, is minimal. I.e., s is a best response if for all strategies r,*

$$C_i^{(s,Q^{-i})} \leq C_i^{(r,Q^{-i})}.$$

We say that protocol Q is in equilibria if $q^{(i)}$ is a best response to Q^{-i} for all agents i.

For one pending agent, the best response for the agent is to transmit deterministically. *I.e.*, for protocols Q in equilibrium, $q_{1,t} = 1$, for all $t \geq 0$. Consequently, $T_{i|1,t} = t + 1$ and $C_{i|1,t} = \Gamma(t+1)$.

Definition 3.4. *A protocol Q is said to be symmetric if $q^{(i)} = q^{(j)}$, for all $i, j \in N$. For symmetric protocols one can use the notation $Q = \langle q \rangle^n$ rather than $Q = \langle q^{(1)}, \ldots, q^{(n)} \rangle$. For the expected cost to an agent we use the notation $C^Q_{k,t}$ instead of $C^Q_{i|k,t}$, as the index i is irrelevant. Likewise, the cost of the protocol can be denoted by C^Q in place of $C^Q_{i|n,0}$.*

For $k \geq 3$, having all the agents continuously transmit (*i.e.*, $q_{k,t} = 1$) is a symmetric, time-independent, protocol in equilibria. Such a protocol is also rather useless as no successful transmissions ever occur[6].

Definition 3.5. *A protocol is called* non-blocking *if for all $k \geq 2$, $t \geq 0$, the transmission probability $q_{k,t} < 1$.*

Note that the expected cost of the game for a time-independent, non-blocking protocol in equilibria is always finite (for the latency cost).

Definition 3.6. *Let $Q = \langle q^{(1)}, \ldots, q^{(n)} \rangle$ be a protocol. Q is said to be efficient if all agents are successful within $D \geq 10n$ time slots, except possibly with exponentially negligible probability ($1/\exp(D)$).*

It does not follow from the definition of efficient protocols that the *expected cost of the game* need be low. Of course, this depends on Γ, but even for latency costs, efficient protocols could have very high latency with some (exponentially small) probability and the expected latency could also be high.

4 Characterization of Symmetric Protocols in Equilibria

In this section we analyze properties of symmetric protocols in equilibria, for general non-negative cost functions. For any symmetric protocol, $Q = \langle q \rangle^n$, where $q = \langle q_{k,t} : 1 \leq k \leq n, 0 \leq t \rangle$, the expected, C, cost for any agent (*e.g.*, Alice) is

$$\begin{aligned} C &= q_{k,t}(1-q_{k,t})^{k-1}\Gamma(t+1) \\ &\quad + (k-1)q_{k,t}(1-q_{k,t})^{k-1}C_{k-1,t+1} \\ &\quad + (1 - kq_{k,t}(1-q_{k,t})^{k-1})C_{k,t+1}. \end{aligned}$$

The first term above is the contribution to the expected cost conditioned on Alice successfully transmitting at time slot t. The second term is the contribution conditioned on some other agent (not Alice) transmitting successfully. The last term is the contribution to the expected cost when there is no successful transmission (either no agent attempts transmission or multiple agents attempt transmission).

[5] For latency costs, classical Markov Decision Theory results [8] show that the best response to a set of time-independent strategies will include some time-independent strategy.

[6] As mentioned in the introduction, for two agents this is not an equilibrium, and $q_{2,t} < 1$.

For an agent strategy q, the strategy $\overline{q}^{(k,t)}$ (respectively, $\underline{q}^{(k,t)}$) is the same as q except that it deterministically transmits (respectively, is quiescent) at time t if $K_t = k$, i.e., $\overline{q}^{(k,t)} = q$ (respectively, $\underline{q}^{(k,t)} = q$) except that $\overline{q}^{(k,t)}_{k,t} = 1$ (respectively, $\underline{q}^{(k,t)}_{k,t} = 0$).

Given that $K_t = k$, the expected cost to Alice, playing strategy $\overline{q}^{(k,t)}$, is

$$C_i^{(\overline{q}^{(k,t)},Q^{-i})} = \alpha_{k,t}\Gamma(t+1) + (1-\alpha_{k,t})C_{k,t+1},$$

where $\alpha_{k,t} = (1-q_{k,t})^{k-1}$ is the probability that none of the other $k-1$ pending agents transmit at time t. Similarly, the expected cost to Alice when playing $\underline{q}^{(k,t)}$ is

$$C_i^{(\underline{q}^{(k,t)},Q^{-i})} = \beta_{k,t}C_{k-1,t+1} + (1-\beta_{k,t})C_{k,t+1},$$

where $\beta_{k,t} = (k-1)q_{k,t}(1-q_{k,t})^{k-2}$ is the probability that exactly one (other) pending agent transmits.

For protocols Q in equilibria, for any k, t such that $0 < q_{k,t} < 1$, it must be that all three strategies, q, $\overline{q}^{k,t}$, and $\underline{q}^{k,t}$, are best responses to Q^{-i}.[7]

We now argue that for symmetric protocols in equilibria, expected cost is monotonically increasing in the number of pending agents and in time.

Lemma 4.1. *Let $Q = \langle q\rangle^n$ be a symmetric protocol in equilibria. For all $k \le n$ and all $t \ge 0$,*

$$C_{k,t} \le C_{k,t+1}, \quad C_{k-1,t} \le C_{k,t}, \quad \text{and } F_{k-1,t} \le F_{k,t}.$$

The following lemma establishes a connection between the transmission probability $q_{k,t}$, and the ratio of future costs $F_{k-1,t+1}/F_{k,t+1}$.

Lemma 4.2. *Let $Q = \langle q\rangle^n$, be a symmetric protocol in equilibrium. For every number of pending agents $1 \le k \le n$ and for every time slot $t \ge 0$, we have*

$$\textit{either}\ \ q_{k,t} = \frac{1}{k-(k-1)\frac{F_{k-1,t+1}}{F_{k,t+1}}}\ ,\textit{or}\ \ q_{k,t} = 1\ .$$

Details on the proofs of Lemmata 4.1 and 4.2 appears in the appendix. The following are immediate consequences of Lemma 4.2.

Corollary 4.1. *If Q is a symmetric protocol in equilibrium then $q_{k,t} \ge \frac{1}{k}$ for every integer $1 \le k$, $0 \le t$.*

Corollary 4.2. *For any constant $0 \le c < 1$, if $F_{k-1,t+1}/F_{k,t+1} < c$, for all $k \le n$, $0 \le t$, then $q_{k,t} = \Theta(\frac{1}{k})$.*

This then implies that the expected latency and the expected maximal latency are both $\Theta(n)$, given that n agents start at time zero.

[7] This follows since both Transmit and being Quiescent are in the support of q at time t with k pending agents.

5 Non-blocking protocols in equilibrium for latency cost

Recall that time-independent protocols are protocols in which $q_{k,t} = q_{k,t'}$ for $t \ne t'$. Thus, for such protocols we can use the notation q_k for transmission probability rather than $q_{k,t}$. We give the following characterization of time-independent, non-blocking protocols, for agents with latency costs.

Theorem 5.1. *There is a unique time-independent, symmetric, non-blocking protocol $\langle q\rangle^n$ in equilibrium for latency cost, $q = \langle q_1, \ldots, q_n\rangle$.*

Furthermore, $q_k \in \Theta(\frac{1}{\sqrt{k}})$, for $1 \le k \le n$.

Proof. Consider agent Alice, one of k pending agents at time t. Assume Alice deviates from q and continuously transmits until successful. This pure strategy is in the support of q, therefore it has an expected cost equal to that of q. As Alice continuously transmits, no agent other than Alice can succeed while Alice is pending. The probability that all agents but Alice are quiescent is fixed at $z = (1-q_k)^{k-1}$. It follows that the expected number of time slots until Alice succeeds is $1/z$, and

$$F_{k,t} = \frac{1}{(1-q_k)^{k-1}}. \tag{5.1}$$

From Lemma 4.2 and equation (5.1) we have,

$$F_{k-1,t} = \frac{1}{(1-q_k)^{k-1}}\left(1 - \frac{1-q_k}{(k-1)q_k}\right). \tag{5.2}$$

Now, substituting $k-1$ for k in equation (5.1), we have that $F_{k-1,t} = \frac{1}{(1-q_{k-1})^{k-2}}$. Substituting this for the left hand side of equation (5.2) we get that

$$\frac{1}{(1-q_{k-1})^{k-2}} = \frac{1}{(1-q_k)^{k-1}}\left(1 - \frac{1-q_k}{(k-1)q_k}\right). \tag{5.3}$$

We first seek to prove inductively that q_k is uniquely determined. For the base case, $k=1$, we have that $q_1 = 1$ and $F_{1,t} = 1$ for all $t \ge 0$. Assume the inductive hypothesis for $k-1$, $k > 1$, then the left hand side of equation (5.3) is some constant. The right hand side of equation (5.3) is a continuous and monotonically increasing function of q_k in the range (0,1) taking values from $-\infty$ to ∞. It follows that q_k is uniquely determined.

We now sketch a proof that $q_k = \Theta(1/\sqrt{k})$, further details in the appendix.

1. It follows from equation (5.1) that $q_k \ge \frac{1}{2\sqrt{k}}$ implies that $F_{k,t} \ge \frac{1}{\left(1-\frac{1}{2\sqrt{k}}\right)^{k-1}}$. From equation (5.2) we learn that $F_{k-1,t} \ge \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-2}}$ implies that

$q_k \geq \frac{1}{2\sqrt{k-1}} \geq \frac{1}{2\sqrt{k}}$. These two derivations prove that $q_k \in \Omega(1/\sqrt{k})$.

2. Similarly, $q_k \leq \frac{1.5}{\sqrt{k}}$ implies that $F_{k,t} \leq \frac{1}{\left(1-\frac{1.5}{\sqrt{k}}\right)^{k-1}}$ and $F_{k-1,t} \leq \frac{1}{\left(1-\frac{1.5}{\sqrt{k-1}}\right)^{k-2}}$ implies that $q_k \leq \frac{1.5}{\sqrt{k}}$.

This then implies that $q_k = O(1/\sqrt{k})$. ∎

6 Efficient Protocols in Equilibria

Given n agents at time zero, Theorem 5.1 implies that the expected latency for (the unique) symmetric, time-independent, protocol in equilibria is exponential in n. Ergo, the probability that even one agent will be successful within any polynomial time bound is exponentially small.

In contrast, efficient protocols ensure that all n agents succeed in linear time except with exponentially small probability. In this section we give a protocol for contention resolution, which is simultaneously efficient, symmetric, and in equilibrium. Obviously, such a protocol *cannot* be time-independent.

6.1 Analysis of Deadline Cost Functions

We now turn aside from latency cost protocols to address strategic behavior under deadline cost functions, such as

$$\Gamma_D(t) = \begin{cases} 0 & \text{For } 0 \leq t < D; \\ 1 & \text{Otherwise.} \end{cases}$$

Or,

$$\Gamma^*_{D,M}(t) = \begin{cases} t & \text{For } 0 \leq t < D; \\ M+t & \text{Otherwise.} \end{cases}$$

Time slot D is referred to as the *deadline*.

The main result of this section is the following theorem that says that for deadlines at least $10n$ time slots away, there exist protocols in equilibria such that all agents succeed before the deadline with high probability:

THEOREM 6.1. *We describe two symmetric protocols, Q_D and $Q^*_{D,M}$.*

1. *Protocol Q_D is in equilibrium for the Γ_D cost function. If the deadline $D > 10n$ then the probability that not all agents succeed prior to the deadline is negligible ($1/e^{\Theta(D)}$).*

2. *Protocol $Q^*_{D,M}$ is in equilibrium for the $\Gamma^*_{D,M}$ cost function. If the deadline $D > 10n$ and $M > \exp(n)$ then the probability that not all agents succeed prior to the deadline is negligible ($1/e^{\Theta(D)}$).*

We first sketch the proof for Γ_D, and then briefly describe the modifications for deadline function $\Gamma^*_{D,\exp(n)}$. Intuitively, as M increases we expect $\Gamma^*_{D,M}$ to become more and more "similar" to Γ_D.

The final protocol for latency cost agents makes them act as if they have deadline costs $\Gamma^*_{D=10n,M=\exp(n)}$. This is achieved as follows: any $k \geq 3$ pending agents at time $t = D$ transmit continuously for M time slots, following which they revert to the time independent protocol for latency cost (see Section 5, and the discussion at the end of this Section). Then, Theorem 6.1 applies and we can conclude that the protocol is not only in equilibrium but also efficient.

6.1.1 The Γ_D deadline cost function.

We begin with the simple case of 2 agents, Alice and Bob. After the deadline expires, all strategies are in equilibrium as they have no effect on the cost.

Assume a symmetric protocol in equilibrium, and consider the transmission probability $q_{2,D-1}$ used by both of the 2 pending agents at time $D-1$. We show that $q_{2,D-1} = 1$. Corollary 4.1 implies that $q_{2,D-1} \geq 1/2 > 0$, assume that $q_{2,D-1} < 1$ — this says that both transmitting and remaining quiescent are in the support of this mixed strategy in equilibrium.

Consider the pure strategy in which Alice chooses to transmit deterministically at time $D-1$. The expected cost to Alice is then equal to the probability that Bob also chooses to transmit, $q_{2,D-1} < 1$. If Alice chooses the pure strategy of remaining quiescent at time $D-1$ then she is doomed to reach the deadline and her cost is exactly 1. *I.e.*, we have unequal expected costs for two pure strategies in the support, contradicting that $q_{2,D-1} < 1$.

It follows that for every symmetric protocol in equilibrium $q_{2,D-1} = 1$. This can be further generalized to whenever k agents are pending at one of the last $k-1$ time slots prior to the deadline, $q_{k,D-k+1} = q_{k,D-k+2} = \cdots = q_{k,D-1} = 1$. One can prove that

LEMMA 6.1. *Consider n agents with deadline cost function Γ_D. Let $\langle q \rangle^n$, be a symmetric protocol in equilibrium for such agents, then, for all $0 \leq t \leq D$ and for any $k > D-t$ we have $q_{k,t} = 1$.*

Likewise, there exists a symmetric protocol in equilibrium for such agents, $\langle q \rangle^n$, such that for all $0 \leq t < D$ for every $k \leq D-t$ we have $q_{k,t} < 1$.

Lemma 6.1 implies that there is some probability $p > 0$ that all k pending agents at time $D-k$ will succeed before the deadline. We remark that given k pending agents at time $D-k$, the probability of even one agent being successful before the deadline D is negligible (super-exponentially small in k). In comparison, Theorem 6.1 says that given k pending agents at time $D-10k$, then all k agents will succeed before D, except with negligible probability.

It is natural to consider the case of $k=2$ pending agents:

LEMMA 6.2. *Symmetric protocols in equilibrium for the Γ_D deadline cost function have*

$$q_{2,t} = \begin{cases} 1/2 & \text{for } 0 \le t \le D-2, \\ 1 & \text{otherwise.} \end{cases}$$

Also, the expected cost $C_{2,t} = (1/2)^{D-t-1}$ for $t \le D-1$.

Proof. For the deadline cost function Γ_D the expected future cost of a single agent at time $t \le D-1$, $F_{1,t} = 0$. However, the expected future cost for one of two pending agents can never be 0 (Since $\langle q \rangle^2$ is a symmetric protocol). It follows from Lemma 4.2 that for $t \le D-2$ we have

$$q_{2,t} = \frac{1}{k-(k-1)\frac{F_{k-1,t+1}}{F_{k,t+1}}} = \frac{1}{2-\frac{0}{F_{2,t+1}}} = 1/2.$$

For deadline cost function Γ_D, the expected cost $C_{2,t}$ equals the probability of remaining unsuccessful until the deadline, $C_{2,t} = (1/2)^{D-t-1}$. ∎

For any deadline D and number of pending agents k we can give a recursive description of the probabilities in equilibrium, $q_{k,t}$, this gives an algorithm for the computation of such $q_{k,t}$, but we now turn to the asymptotic analysis of such equilibria. Obviously, for all $0 \le t < D$, $2 \le k \le n$, either $F_{k,t} \le 2F_{k-1,t}$ or $F_{k,t} > 2F_{k-1,t}$. In the latter case, it follows from Lemma 4.2 that

$$q_{k,t-1} = \frac{1}{k-(k-1)\frac{F_{k-1,t}}{F_{k,t}}} < \frac{1}{k-(k-1)/2} < 2/k.$$

We now describe a rooted tree, $T=(V,E)$, with weights on the edges. For edge z, $w(z)$ is the weight of the edge. Vertices $v \in V$ have labels $\ell(v) = (k,t)$ for some $1 \le k \le n$, $0 \le t \le D$. Not all the $n(D+1)$ possible labels need appear on some vertex $v \in V$, and the same label may appear multiple times ($\ell(v) = \ell(v')$, $v \ne v'$). The root vertex r is assigned the label $\ell(r) = (n,0)$, and is the only vertex so labelled.

Given $v \in V$, with $\ell(v) = (k,t)$, $2 \le k \le n$, $0 \le t < D$, we attach descendants to v as follows:

- If $F_{k,t} \le 2F_{k-1,t}$, then v has one descendant, x, with $\ell(x) = (k-1,t)$. Edge (v,x) is given weight $w(v,x) = F_{k,t}/F_{k-1,t}$, note that $w(v,x) \le 2$. Such edges, where v has a single descendant, are called *doubling edges* and the set of all such edges is denoted by E_d.
- Otherwise, v has two descendants, y, and x, where $\ell(y) = (k-1,t+1)$ and $\ell(x) = (k,t+1)$. The weight $w(v,y) = \beta_{k,t} = (k-1)q_{k,t}(1-q_{k,t})^{k-2}$, and $w(v,x) = 1-\beta_{k,t}$. (Note that $\beta_{k,t}$ is defined as in Section 4.) What this means is that should agent Alice, one of k pending agents at time t, choose to be quiescent at time t, then with probability $\beta_{k,t}$, one of the agents other than Alice will be successful at time t. The weight $w(v,x) = 1-\beta_{k,t}$ and represents the probability that in the same setting, no agent will be successful and the time slot will have been "wasted".

Let L_0 be the set of vertices $v \in V$ with labels $\ell(v) = (1,t)$, $0 \le t < D$. Let L_1 be the set of vertices $v \in V$ with labels $\ell(v) = (k,D)$, $1 \le k \le n$. The set $L_0 \cup L_1$ is exactly the set of leaves in T.

For any leaf v, where $\ell(v) = (k,t)$, we define the real value $c(v) = C_{k,t}$. *I.e.*, for $v \in L_0$, $c(v) = 0$ and for $v \in L_1$, $c(v) = 1$. An internal vertex v with two descendants, x,y, has $c(v) = w(v,x)c(x) + w(v,y)c(y)$, an internal vertex v with one descendant, x, has $c(v) = w(v,x)c(x)$. It follows from the recursive construction of T and from the recursive evaluation of $c(\cdot)$ that $c(r) = C_{n,0}$.

For a leaf v in T let $P(v)$ denote the set of edges along the path from the root r to v. One can rearrange the recursive summation for the value of $c(r)$ as follows:

$$\begin{aligned} c(r) &= \sum_{v \in L_0} \left(c(v) \prod_{(u,v) \in P(v)} w(u,v) \right) \\ &\quad + \sum_{v \in L_1} \left(c(v) \prod_{(u,v) \in P(v)} w(u,v) \right) = \\ &\quad \sum_{v \in L_1} \left(\prod_{(u,v) \in P'(v)} w(u,v) \prod_{(u,v) \in P''(v)} w(u,v) \right), \end{aligned}$$

where $P'(v) = P(v) \cap E_d$ and $P''(v) = P(v) \cap (E - E_d)$.

LEMMA 6.3. *For all edges $(u,v) \in E - E_d$, $w(u,v) < 1 - \frac{2}{e^2} = \beta < 1$.*

The size of L_1 is no more than $\sum_{k=1}^{n} \binom{k+D}{D} \le n\binom{n+D}{D}$. The product of edges weights for edges in E_d is at most 2^n, since there are at most n such edges (this follows since a doubling edge decreases the first coordinate of the label by one, and vertices with labels $(1,t)$ are leaves). The product of edge weights for edges in $E - E_d$ along some path from r to $v \in L_1$ decreases exponentially with the path length, which is at least D.

It follows that

$$\prod_{(u,v) \in P(v) \cap (E-E_d)} w(u,v) \le \beta^D, \text{ and thus}$$

$$c(r) \le n \binom{n+D}{D} 2^n \beta^D.$$

For $D = bn$, $c(r) \le e^{\Theta(n \ln b) - nb \ln(1/\beta)}$ which is exponentially decreasing in $D = bn$. The value $c(r)$ is the probability that a specific agent will fail to successfully transmit before the deadline. It follows that the probability that all n agents are successful prior to the deadline is at least $1 - nC_{n,0}$.

6.1.2 The $\Gamma^*_{D,M}$ deadline cost function. To generalize this argument for $\Gamma^*_{D,M}$ we define the protocol to use the time independent latency cost protocol after the deadline has passed. This protocol implies a future cost $F_{k,D+1}$ of $\exp(\sqrt{k}) < \exp(\sqrt{n})$. Choose M to be $\exp(n)$ and define vertices v labelled $(k, D+1)$ to be leaves with $c(v) = M + \exp(\sqrt{n})$ (analogously to the set L_1 above). We get that $C_{n,0} \le c(r)$, and an argument similar to the above goes through.

7 Discussion and Open Problems

Our work and extensions thereof can be cast within extensions of the framework of multicommodity flow and Waldrop equilibria. Multicommidity flow in general, and much of the previous work on strategic behavior for multicommidity flow, implicitly assumes a steady state where the (s_i, t_i) flow for agent i has an associated "flow rate", the number of gallons per minute or the bandwidth.

Our study above focuses on another important parameter which is the flow duration, *i.e.*, agent i requires flow from s_i to t_i for a finite duration. *E.g.*, the flow width is the bandwidth required to transmit MPEG, the flow duration is the length of the movie. One may consider "congestion" as a special case or variant of the edge latency functions used in Waldrop equilibria.

Infinite duration flow (steady state flow) fails to capture fundamental problems and common sense solutions such as leaving home at 10:00 AM so as to avoid rush hour traffic.

Generalizing our approach used above for single link single packet congestion control so as to deal with more general finite duration multicommidity flow problems is an intriguing direction for future research. Even more generally, the analysis above shows that strategic behavior may lead to some form of "altruism". The key intuition is that "altruism" arises — not on moral grounds — but because it is simply another way of removing the competition from play. One can imagine describing altruistic behavior in various file sharing networks similarly. Understanding "altruism" is itself a major research goal.

Yet other directions of future research, include the following:

1. Deal with the case where the number of agents pending is unknown. This is probably the single most important open problem.

2. Extend the results above to deal with packets that are generated over time. It seems that appropriate modifications of the above arguments should work for at least some packet generation rates.

3. Extend these results to other congestion functions, not necessarily only a $0/\infty$ step function. Many alternatives are possible, perhaps most interesting and well motivated is allowing some probabilistic result when contention occurs. As a motivating example consider the problem of cache overflow, a critical threshold is reached when pages used repeatedly by a set of active tasks fill the cache. Performance may drop by many orders of magnitude.

4. Dealing with yet other agent utilities, not only deadlines and linear latency. One can imagine a host of other latency based utility functions, Γ, each of which can find justification in a variety of different scenarios. One can also imagine that different agents have different utilities.

References

[1] Eitan Altman, Rachid El Azouzi, and Tania Jiménez. Slotted aloha as a game with partial information. *Computer Networks*, 45(6):701–713, 2004.

[2] Eitan Altman, Dhiman Barman, Rachid El Azouzi, and Tania Jiménez. A game theoretic approach for delay minimization in slotted aloha.

[3] Feigenbaum and Shenker. Distributed algorithmic mechanism design: Recent results and future directions. *BEATCS: Bulletin of the European Association for Theoretical Computer Science*, 79, 2003.

[4] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A bgp-based mechanism for lowest-cost routing, 2002.

[5] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Lecture Notes in Computer Science*, 1563:404–413, 1999.

[6] Allen MacKenzie and Stephen B. Wicker. Stability of multipacket slotted aloha with selfish users and perfect information. In *INFOCOM*, 2003.

[7] Martin J. Osborne and Ariel Rubenstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.

[8] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1994.

[9] Raphael Rom and Moshe Sidi. *Multiple access protocols: performance and analysis.* Springer-Verlag New York, Inc., New York, NY, USA, 1990.

[10] Tim Roughgarden. *Selfish Routing And The Price Of Anarchy.* The MIT Press, Cambridge, Massachusetts, 2005.

A Proofs from Section 4

Proof of Lemma 4.1: Consider the case of k pending agents, one of which is agent i. Recall that in equilibrium, each agent has the same expected cost for playing either q or $\overline{q}^{(k,t)}$ or $\underline{q}^{(k,t)}$. Namely,

$$\begin{aligned} C_{k,t} &= \alpha_{k,t}\Gamma(t+1) + (1-\alpha_{k,t})C_{k,t+1} \\ &= \beta_{k,t}C_{k-1,t+1} + (1-\beta_{k,t})C_{k,t+1}, \end{aligned}$$

where $\alpha_{k,t}$ denotes the probability that none of the other $k-1$ pending agents transmit at time t, and $\beta_{k,t}$ denotes the probability that exactly one of the other $k-1$ pending agents transmits at time t.

All agents pending at time $t+1$ will not have success prior to time $t+2$ and thus have cost $\geq \Gamma[t+2]$, i.e.,

$$C_{k,t+1} \geq \Gamma[t+2] \geq \Gamma[t+1] \ .$$

This implies that,

$$C_{k,t} = \alpha_{k,t}\Gamma[t+1] + (1-\alpha_{k,t})C_{k,t+1} \leq C_{k,t+1} \ ,$$

establishing that $C_{k,t} \leq C_{k,t+1}$.

By the fact that $C_{k,t} \leq C_{k,t+1}$, we get that

$$C_{k,t} \geq \beta_{k,t}C_{k-1,t+1} + (1-\beta_{k,t})C_{k,t}$$

Hence,

$$\beta_{k,t}C_{k,t} \geq \beta_{k,t}C_{k-1,t+1} \geq \beta_{k,t}C_{k-1,t} \ ,$$

where the second inequality follows since $C_{k-1,t+1} \geq C_{k-1,t}$. This establishes that $C_{k-1,t} \leq C_{k,t}$.

By definition of $F_{k,t}$:

$$F_{k,t} = C_{k,t} - \Gamma(t) \geq C_{k-1,t} - \Gamma(t) = F_{k-1,t} \ ,$$

showing that $F_{k-1,t} \leq F_{k,t}$. ∎

Proof of Lemma 4.2: Consider Alice, one of k pending agents at time t. All agents but Alice follow strategy q. For symmetric protocols, transmission probabilities $q_{k,t}$ are always > 0. If $q_{k,t} = 1$ then we are done. Otherwise, in equilibrium, all agents, including Alice, have the same expected cost whether playing q, $\overline{q}^{(k,t)}$, $\underline{q}^{(k,t)}$. I.e.,

$$\begin{aligned} C_{k,t} &= \alpha_{k,t}\Gamma(t+1) + (1-\alpha_{k,t})C_{k,t+1} \\ &= \beta_{k,t}C_{k-1,t+1} + (1-\beta_{k,t})C_{k,t+1} \end{aligned}$$

Therefore,

$$\alpha_{k,t}(\Gamma(t+1) - C_{k,t+1}) = \beta_{k,t}(C_{k-1,t+1} - C_{k,t+1}).$$

Substituting $\alpha_{k,t} = (1-q_{k,t})^{k-1}$ and $\beta_{k,t} = k-1)q_{k,t}(1-q_{k,t})^{k-2}$, we get that

$$\begin{aligned} &(1-q_{k,t})^{k-1}(\Gamma(t+1) - C_{k,t+1}) = \\ &\qquad (k-1)q_{k,t}(1-q_{k,t})^{k-2}(C_{k-1,t+1} - C_{k,t+1}). \end{aligned}$$

Since by assumption $q_{k,t} \neq 1$, dividing by $(1-q_{k,t})^{k-2}$ results in

$$\begin{aligned} &(1-q_{k,t})(\Gamma(t+1) - C_{k,t+1}) = \\ &\qquad (k-1)q_{k,t}(C_{k-1,t+1} - \Gamma(t+1) + \Gamma(t+1) - C_{k,t+1}). \end{aligned}$$

As $F_{k,t+1} = C_{k,t+1} - \Gamma(t+1)$ it follows that

$$(1-q_{k,t})F_{k,t+1} = (k-1)q_{k,t}(F_{k,t+1} - F_{k-1,t+1}),$$

and the claim follows. ∎

Remark A.1. *In the proof of Lemma 4.2, we implicitly assume that $C_{k,t}$ is finite. This holds for any non-blocking protocol with respect to the latency cost function.*

B Proof of Theorem 5.1

In this appendix we give the proof of the claims that establish Theorem 5.1.

Proposition B.1. *If $q_k \geq \frac{1}{2\sqrt{k}}$ then $F_{k,t} \geq \frac{1}{1-\frac{1}{2\sqrt{k}}}^{k-1}$.*

Proof. The claim follows directly from equation (5.1), since $F_{k,t}$ is monotonically increasing in q_k. ∎

Next we provide a lower bound on q_k as a function of $F_{k,t}$.

Proposition B.2. *If $F_{k-1,t} \geq \frac{1}{1-\frac{1}{2\sqrt{k-1}}}^{k-2}$ then $q_k \geq \frac{1}{2\sqrt{k-1}} \geq \frac{1}{2\sqrt{k}}$.*

Proof. Fix k and investigate the following rational function (that stems from equation (5.2)).

$$L(x) = \frac{1}{(1-x)^{k-1}}\left[1 - \frac{1-x}{(k-1)x}\right].$$

The function $L(\cdot)$ is monotonically increasing in the range $[0,1]$. To see this note that the derivative of $L(x)$ is

$$\begin{aligned} L'(x) = \frac{k-1}{(1-x)^k} &+ \frac{k-2}{(1-x)^{k-1}}\frac{1}{(k-1)x} \\ &+ \frac{1}{(1-x)^{k-2}}\frac{1}{(k-1)x^2}. \end{aligned}$$

Consider the rational function

$$R(x) = L(x) - \frac{1}{\left(1 - \frac{1}{2\sqrt{k-1}}\right)^{k-2}}.$$

Since $R(x)$ is monotonically increasing in the range $[0, 1]$, it has a single root, z, in $[0, 1]$. We show that $R(\frac{1}{\sqrt{k-1}}) > 0$ and $R(\frac{1}{2\sqrt{k-1}}) < 0$, therefore $z \in [\frac{1}{2\sqrt{k-1}}, \frac{1}{\sqrt{k-1}}]$. Thus,

$$R(x) > 0, x \in [0, 1] \implies x > z > \frac{1}{2\sqrt{k-1}}.$$

First we show that $R(\frac{1}{\sqrt{k-1}}) > 0$:

$$\begin{aligned}
R(\frac{1}{\sqrt{k-1}}) &= \frac{1}{(1-\frac{1}{\sqrt{k-1}})^{k-1}}\left[1 - \frac{1-\frac{1}{\sqrt{k-1}}}{(k-1)\frac{1}{\sqrt{k-1}}}\right] \\
&- \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-2}} \\
&= \frac{1}{\left(1-\frac{1}{\sqrt{k-1}}\right)^{k-1}}\left[1 - \frac{1}{\sqrt{k-1}} + \frac{1}{k-1}\right] \\
&- \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-2}} \\
&> 0.
\end{aligned}$$

Next we show that $R(\frac{1}{2\sqrt{k-1}}) < 0$.

$$\begin{aligned}
R(\frac{1}{2\sqrt{k-1}}) &= \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-1}}\left[1 - \frac{1-\frac{1}{2\sqrt{k-1}}}{(k-1)\frac{1}{2\sqrt{k-1}}}\right] \\
&- \frac{1}{(1-\frac{1}{2\sqrt{k-1}})^{k-2}} \\
&= \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-1}}\left[1 - \frac{1-\frac{1}{2\sqrt{k-1}}}{\sqrt{k-1}/2}\right] \\
&- \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-2}} \\
&= \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-1}}\left[1 - \frac{2}{\sqrt{k-1}} + \frac{1}{k-1}\right] \\
&- \frac{1}{(1-\frac{1}{2\sqrt{k-1}})^{k-2}} \\
&\leq \frac{1}{\left(1-\frac{1}{2\sqrt{k-1}}\right)^{k-1}}\left[-\frac{1.5}{\sqrt{k-1}} + \frac{1}{k-1}\right] \\
&< 0.
\end{aligned}$$

■

C Proofs from Section 6

Proof of Lemma 6.1: The proof is by double induction, the outer induction is on the number of agents, and the inner induction is on the number of time slots remaining until the deadline.

We've previously argued that $q_{2,D-1} = 1$ but the same argument also shows that $q_{k,D-1} = 1$ for all $k \leq n$.

The induction hypothesis for the outer induction is that for some $k < n$, and all $t > D - k$, $q_{k,t} = 1$. The inductive step is to prove that for all $t > D - (k+1)$, $q_{k+1,t} = 1$. The base case for the outer induction is that $q_{2,D-1} = 1$.

To prove the outer induction for $k+1$, we use an inner induction on the number of time slots remaining until the deadline. The base for the inner induction is that $q_{k+1,D-1} = 1$. The inner induction hypothesis is that $q_{k+1,t} = 1$ for some $t > D - (k+1) + 1 = D - k$. Let Alice be one of these $k+1$ pending agents that plays Quiescent at time $t-1$. Even if some agent other than Alice is successful at time $t-1$, there will still be $\geq k$ pending agents (including Alice) at time t. By the outer induction, Alice is doomed not to succeed before the deadline, and thus $q_{k+1,t-1} < 1$ cannot be in equilibrium. ■

Proof of Lemma 6.3: A vertex labelled (k, t) that has two descendants implies that $q_{k,t} < \frac{2}{k}$. The term $\beta_{k,t} = (k-1)q_{k,t}(1-q_{k,t})^{k-2}$ as a function of $q_{k,t}$ is monotonically decreasing in the range $[\frac{1}{k}, 1]$. Hence,

$$\begin{aligned}
\frac{2}{e^2} &< (k-1)\frac{2}{k}\left(1-\frac{2}{k}\right)^{k-2} < \beta_{k,t} \\
&< (k-1)\frac{1}{k}\left(1-\frac{1}{k}\right)^{k-2} < \frac{1}{2}\ .
\end{aligned}$$

■