Efficient Contention Resolution Protocols for Selfish Agents

Amos Fiat Yishay Mansour Uri Nadav

August 15, 2007

"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 [1].

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 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[2] 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 commons", 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.

[1] More precisely, this protocol is the optimal history independent protocol, however, there exists protocols that achieve better throughput based on the channel's history. As the $1/k$ transmission probability attains expected linear delay, we use it as a benchmark and regard it as the socially optimal protocol.

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

We call such an equilibrium 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 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[3]. Additionally, we deal with symmetric protocols, where all agents follow the same set of rules[4]. Our primary interest is in

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

[4]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).

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 $e^{\Theta(\sqrt{n})}$ (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 $15n$ 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 studies protocols for agents with deadline cost functions (as well as agents with latency costs).

We seek equilibria[5] 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

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

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 Wardrop equilibria for traffic, cast 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 truthfulness, 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 $i \in \{1, \dots, n\}$ 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. 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 $t \geq D$). Section 7 deals with exponentially increasing/ decreasing marginal cost functions *i.e.*, $\Gamma(t) = \sum_{j=0}^{t} \delta^j$, for some constant δ.

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 $0 \le t, t'$. 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*[6].

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, $1 \le i \le n$ has strategy $q^{(i)}$.*

Fix a protocol Q. We define the expected cost of the protocol for agent i, $C_i^Q = \mathrm{E}[\Gamma(T_i)]$. 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 $C^Q_{i|k,t} = \mathrm{E}[\Gamma(T_{i|k,t})]$, and define the expected future cost $F^Q_{i|k,t} = C^Q_{i|k,t} - \Gamma(t) = \mathrm{E}[\Gamma(T_{i|k,t})] - \Gamma(t)$. (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 $j \ne i$ 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 $j \ne i$ use $q^{(j)}$, is minimal. I.e., s is a best response if for all strategies r,*

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

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 $t \ge 0$. 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 $k \ge 3$, 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[7].

Definition 3.5. *A protocol is called* non-blocking *if for all $k \ge 2$, $t \ge 0$, 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 $D = O(n)$ time slots, except possibly with exponentially negligible probability $(1/\exp(n))$.*

[6] 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.

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

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).

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

$$\alpha_{k,t} = (1-q_{k,t})^{k-1} , \tag{4.1}$$

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

$$\beta_{k,t} = (k-1)q_{k,t}(1-q_{k,t})^{k-2} \tag{4.2}$$

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}.[8]

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

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

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

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

Proof. 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 $\alpha_{k,t}, \beta_{k,t}$ are as defined in Equations 4.1 and 4.2 respectively.

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

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

This implies that,

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

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

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

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

Hence,

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

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

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

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

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

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*

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

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 strictly positive. If $q_{k,t} = 1$ then we are done. Otherwise, in equilibrium, all agents, including Alice, have the same expected cost whether playing q, $\overline{q}^{(k,t)}$ or $\underline{q}^{(k,t)}$. *I.e.*,

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

Expression	Definition
$\Gamma(t)$	The cost for an agent who successfully transmits at time t
$\gamma(t)$	The marginal cost added after staying in the system at time t
$q_{k,t}$	Transmission probability when k agents are pending at time t
$\alpha_{k,t}$	The probability that none of the other $k-1$ pending agents transmit at time t
$\beta_{k,t}$	The probability that exactly one (other) pending agent transmits at time t
$C_{k,t}$	The expected cost given k agents are pending at time t.
$F_{k,t}$	The expected future cost given k agents are pending at time t

Table 1: Notations in use.

Therefore,

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

Substituting according to Equations 4.1,4.2, we get that

$$(1-q_{k,t})^{k-1}(\Gamma(t+1) - C_{k,t+1}) = (k-1)q_{k,t}(1-q_{k,t})^{k-2}(C_{k-1,t+1} - C_{k,t+1}).$$

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

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

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 4.3. *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.*

The following are immediate consequences of Lemma 4.2.

Corollary 4.4. *If Q is a symmetric protocol in equilibrium then $q_{k,t} \geq \frac{1}{k}$ for every integer $1 \leq k$, $0 \leq t$.*

Corollary 4.5. *For any constant $0 \leq c < 1$, if $F_{k-1,t+1}/F_{k,t+1} < c$, for all $k \leq n$, $0 \leq t$, 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.

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 $t \neq t'$. Thus, for such protocols we can use the notation q_k for transmission probability rather than $q_{k,t}$. When time-independent strategies are used for the latency cost function, the future cost depends only on the number of pending agents, *i.e.*, $F_{k,t} = F_{k,t'}$ for $t \neq t'$. We can therefore use the notation F_k for future cost rather than $F_{k,t}$. 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 $1 \leq k \leq n$.*

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

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

From Lemma 4.2 and equation (5.1) we have,

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

Now, substituting $k-1$ for k in equation (5.1), we have that $F_{k-1} = \frac{1}{(1-q_{k-1})^{k-2}}$. 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 $t \geq 0$. 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 are found in Appendix B.

1. It follows from equation (5.1) that $q_k \geq \frac{1}{\sqrt{k}+1}$ implies that $F_k \geq \frac{1}{\left(1-\frac{1}{\sqrt{k}+1}\right)^{k-1}}$. From equation (5.2) we learn that $F_{k-1} \geq \frac{1}{\left(1-\frac{1}{\sqrt{k-1}+1}\right)^{k-2}}$ implies that $q_k \geq \frac{1}{\sqrt{k-1}+1} \geq \frac{1}{\sqrt{k}+1}$. These two derivations prove that $q_k \geq \frac{1}{\sqrt{k}+1}$.
2. Similarly, $q_k \leq \frac{2}{\sqrt{k}}$ implies that $F_k \leq \frac{1}{\left(1-\frac{2}{\sqrt{k}}\right)^{k-1}}$ and $F_{k-1} \leq \frac{1}{\left(1-\frac{2}{\sqrt{k-1}}\right)^{k-2}}$ implies that $q_k \leq \frac{2}{\sqrt{k}}$.

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

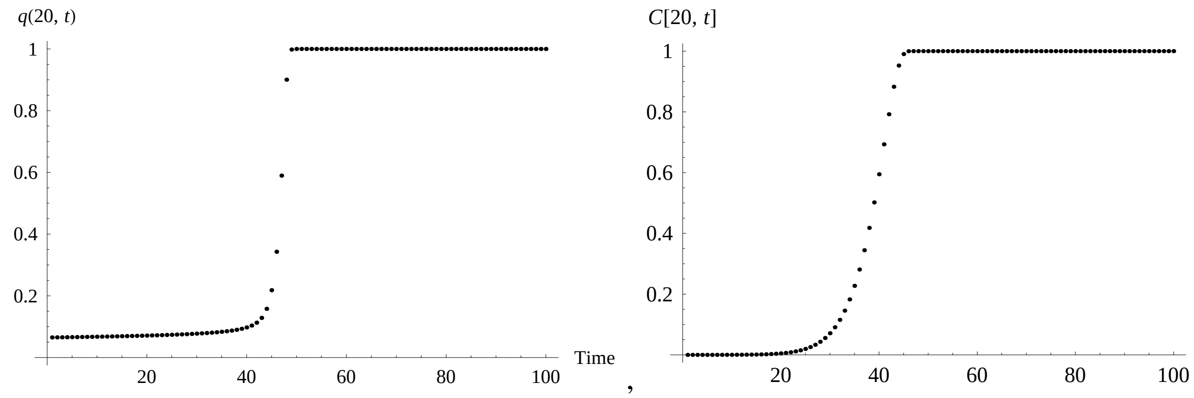

Figure 1: Numerical solution of the deadline equilibrium equations with deadline set at the 100 time slot. The number of agents is fixed to $n = 20$. On the left, the transmission probability of fixed n agents as a function of the time is described. On the right, the probability of reaching the deadline before a successful transmission as a function of the time.

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.

To achieve efficient protocols for agents with latency costs, we turn aside from latency cost protocols to address strategic behavior under deadline cost functions.

Figure 1 shows the numerical solution of the symmetric protocol in equilibrium, for $n = 20$ agents and a deadline at time $D = 100$. There seems to be phase transition of the transmission probabilities approximately at time $D - 3.5n$. Similar phenomena appear for other values of n. Unfortunately, we are unable to prove this empirical observation. What we can show is that for time slots close to the deadline (say $D - (1 + \epsilon)n$), the transmission probability is arbitrarily close to one. We also show that the probability that even a single agent is delayed until the deadline is negligible if the deadline is at least $15n$ time slots away. This implies that agents transmit with probability $O(1/k)$ during a large fraction of the time slots $t \in \{D - 15n, \dots, D\}$.

We consider two related deadline cost functions, the first is a pure deadline, the 2nd is a combination of latency costs plus a deadline.

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

Or,

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

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 $15n$ 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 $D > 15n$ 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 $D > 15n$ 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.

Our efficient protocol for latency cost agents "tricks" them into acting as if they have deadline costs $\Gamma^*_{D=15n,M=\exp(n)}$. This is achieved as follows: any $k \ge 3$ 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.

Until further notice, we now restrict attention to 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 4.4 implies that $q_{2,D-1} \ge 1/2 > 0$. 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 6.2. *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 $0 \le t \le D$ 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 $0 \le t < D$ for every $k \le D - t$ we have $q_{k,t} < 1$.

Proof. 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. ∎

Lemma 6.2 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 $D - 15k$, 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 6.3. *Symmetric protocols in equilibrium for the Γ_D deadline cost function have*

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

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

Proof. For the deadline cost function Γ_D the expected future cost of a single agent at time $t \leq D-1$, $F_{1,t} = 0$. 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 $t \leq D-2$ 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 $0 \leq t < D$, $2 \leq k \leq n$, either $F_{k,t} \leq 2F_{k-1,t}$ or $F_{k,t} > 2F_{k-1,t}$. In the latter case, it follows from Lemma 4.2 that

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

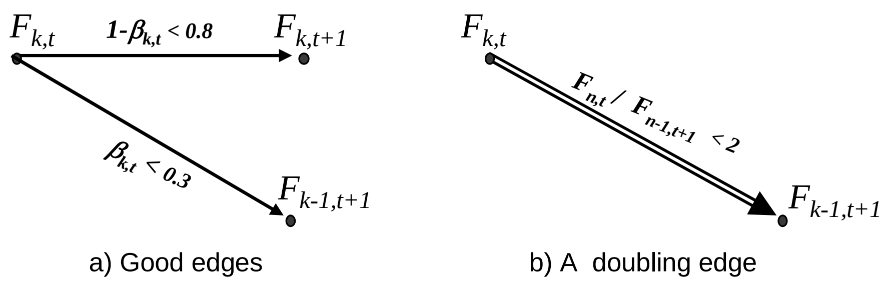

Figure 2: a) Good edges; weight strictly less than 1. b) A doubling edge; weight at most 2.

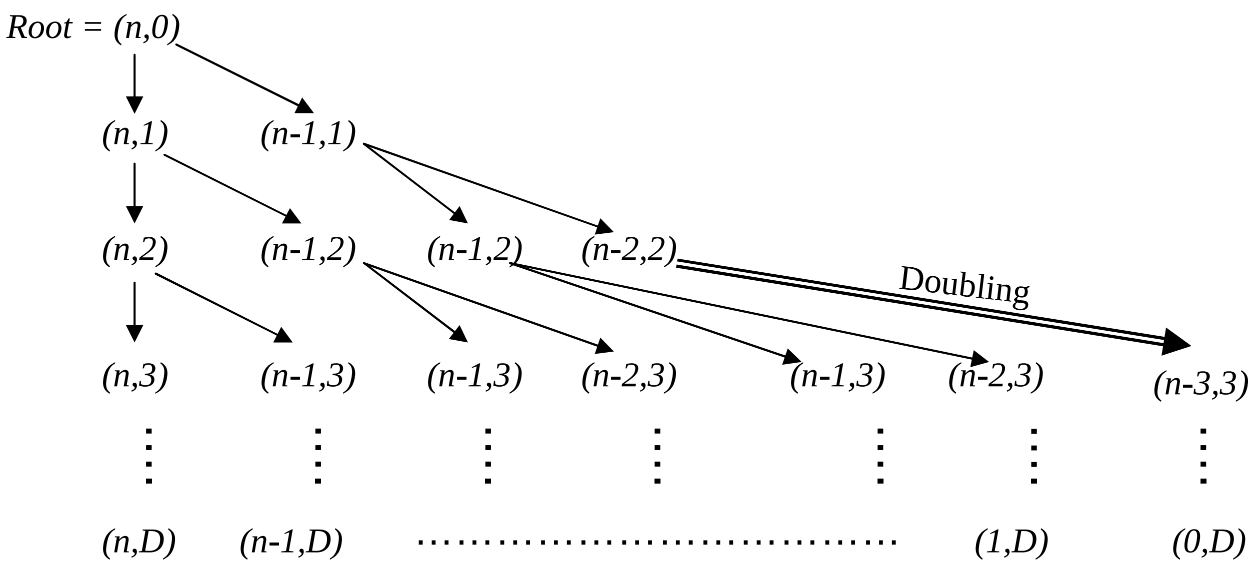

Figure 3: An example of a tree formed as described in the analysis of Theorem 6.1. Notice that multiple vertices may share the same label.

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 (as illustrated in Figure 3). 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 labeled.

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 (as illustrated in Figure 2):

- If $F_{k,t+1} \le 2F_{k-1,t+1}$, then v has one descendant, x, with $\ell(x) = (k-1, t+1)$. Edge (v, x) is given weight $w(v, x) = F_{k,t}/F_{k-1,t+1}$. Note that $w(v, x) \le 2$, since $F_{k,t} \le F_{k,t+1} \le 2F_{k-1,t+1}$, where the first inequality follows Lemma 4.1. 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 $w(v, y) = \beta_{k,t}$ (See Equation 4.2 to recall the definition of β_{kt}), and $w(v, x) = 1 - \beta_{k,t}$.

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:

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

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

Lemma 6.4. *For all edges $g \in E - E_d$, $w(g) < 1 - \frac{2}{e^2} = \beta < 1$.*

Proof. A vertex labeled (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 $[\frac{1}{k-1}, 1]$. Hence, for $k \ge 2$,

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

∎

The size of L_1 is no more than $\sum_{k=1}^{n} \binom{D}{k}$ which is less than $\leq n\binom{D}{n}$ for a deadline $D \geq 2n$. 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 $D - n$.

It follows that

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

$$c(r) \leq n\binom{D}{n} 2^n \beta^{D-n}.$$

For $D = bn$, $\binom{D}{n} < (eb)^n$ and therefore $c(r) \leq e^{\Theta(n \ln b) - nb \ln(1/\beta)}$ 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 $1 - nC_{n,0}$. Setting $b = 15$ suffices to have the cost $c(r)$ diminishes in n when the deadline is set to bn.

Remark 6.5. *The choice of the constant 2 in the analysis (i.e., vertex (k, t) has a single descendant if $F_{k,t+1} \leq 2F_{k-1,t+1}$ and two descendants otherwise) was arbitrary. Optimizing, we can reduce the requirement on the deadline in Theorem 6.1 from being at least $15n$ to being at least $12.3n$. This is done by changing the construction rule so that a vertex (k, t) has a single descendant if $F_{k,t+1} < 3.4F_{k-1,t+1}$ and two descendants otherwise.*

We now turn to the $\Gamma^*_{D,M}$ cost function, a combination of both latency and deadline costs.

This cost function arises in the design of an efficient protocol for latency costs. The benign behavior of selfish agents when faced with a deadline as shown above, suggests a an efficient time-dependent protocol for latency cost:

Given n agents with latency costs, define the deadline $D = 15n$, the time dependent protocol we suggest is as follows:

1. Follow the strategic behavior of a non-blocking Nash equilibria up to time D, while keeping in mind that

2. at time D - if more than 3 agents remain then all agents transmit for 2^n time steps.

Effectively, this means that the utility function for the agents is no longer latency costs but rather $\Gamma^*_{D,M}$, where $M \geq \exp(n)$.

This protocol implies a future cost for $k \geq 3$ agents at time D, $F_{k,D}$, to be $\geq \exp(n)$. Define vertices v labelled (k, D), $k \geq 3$, to be leaves of a tree with $c(v) \geq M$ (analogously to the set L_1 above). Define vertices v labelled $(1, j)$, $j \leq D$, to be leaves with $c(v) = j$. Likewise, vertices $(1, D)$ and $(2, D)$ are leaves with $c(v) = O(1)$. Define the weights of the edges as above, by construction, $C_{n,0} = c(r)$. Solving for the $q_{k,t}$ as above gives us that the probability that any agent remain in the system at time D is negligible.

It is important to remark is that the expected cost is not linear, rather - the protocol completes in linear time with all but negligible probability.

7 Exponentially Growing and Shrinking Marginal Costs

In Section 5 we analyzed the performance of the unique symmetric, time-independent protocol in equilibrium for the linear latency cost function. In this section we show that such a symmetric equilibrium in time-independent strategies exists for a broader class of cost functions of the form $\Gamma(t) = \sum_{j=0}^{t} \delta^j$. When $\delta < 1$, such a cost function may be regarded as a discount factor, common in the analysis of repeated games. For such cost functions it is convenient to consider a *rent* function $\gamma(t)$ for the marginal cost at time slot t, *i.e.*, $\gamma(t) = \delta^t$.

It turns out that $\delta = 1$ is a critical point for the expected time it takes to resolve contention in symmetric equilibrium. When $\delta < 1$ the transmission probability in equilibrium is bounded from below by a constant that depends only on δ and therefore the expected time for transmission is exponential in the number of agents. When $\delta > 1$ the transmission probability is $\Theta(1/k)$ and the expected time for successful transmission is linear in the number of agents. Recall that Theorem 5.1 shows that time independent protocols for linear costs (when $\delta = 1$) have expected latency $\Theta(\exp(\sqrt{n}))$.

In terms of costs, however, the situation is quite the opposite. When $\delta < 1$ there is an upper bound of $1/(1-\delta)$ on the cost. For $\delta > 1$, even in the case of optimal scheduling of transmitted packets, the expected time remains linear in n and therefore the expected cost is exponential.

Theorem 7.1 characterizes the transmission probability in equilibrium for a whole range of $0 < \delta < e/(e-1)$. However, for $\delta \geq e/(e-1)$, it can be shown that no time independent symmetric protocol attains finite expected cost. As our criterion for comparing actions relies solely on the expected cost, we do not consider exponential rent functions with $\delta > e/(e-1)$.

Theorem 7.1. *There is a unique time-independent, symmetric, non-blocking protocol $\langle q \rangle^n$ in equilibrium $q = \langle q_1, \ldots, q_n \rangle$ for rent functions of the form δ^t Furthermore,*

- *For $\delta < 1$, the transmission probability $q_k < 1 - \delta,\ 2 \leq k \leq n$.*
- *For $\delta = 1$, the transmission probability $q_k \in \Theta(\frac{1}{\sqrt{k}})$, for $1 \leq k \leq n$ (This is Theorem 5.1).*
- *For $1 < \delta < e/(e-1)$, the transmission probability $q_k < \frac{1}{k-1} \cdot \frac{\delta}{e(\delta-1)},\ 2 \leq k \leq n$.*

For the proof of Theorem 7.1 we use the fact that time is irrelevant in the description of equilibrium in time-independent strategies, for such cost functions, as shown in Lemma 7.2.

Lemma 7.2. *For a symmetric time-independent strategy $q = \langle q_1, \ldots, q_n \rangle$, and a rent function $\gamma(t) = \delta^t$,*

$$C_{k,t+1} = 1 + \delta C_{k,t}\ ;\ F_{k,t+1} = \delta F_{k,t}$$

Proof. Consider a rent function δ^t, $\delta < 1$. First, observe that for an exponential rent function,

$$\Gamma(t+1) = \sum_{j=0}^{t+1} \delta^t = 1 + \sum_{j=1}^{t+1} \delta^t = 1 + \delta\Gamma(t)\ .$$

Also, at every slot, the expected time until successful transmission T, for some pending agent Alice, depends only on the number of pending agents. Recall that the random variable $T_{k,t}$ describes the time of a successful transmission when there are k agents pending at time t. For time-independent strategies, the probability of a successful transmission in i time slots depends only the number of pending agents, *i.e.*,

$$\Pr[T_{k,t} = t + i] = \Pr[T_{k,s} = s + i] \ .$$

This means that for every i,k,t, the random variable $T_{k,t+i}$ has the same distribution as the random variable $(T_{k,t} + i)$,

$$\Pr[(T_{k,t}+i) = t+M] = \Pr[T_{k,t} = t+M-i] = \Pr[T_{k,t+i} = t+i+M-i] = \Pr[T_{k,t+i} = t+M] \ ,$$

and therefore,

$$C_{k,t+1} = \mathrm{E}[\Gamma(T_{k,t+1})] = \mathrm{E}[\Gamma(T_{k,t} + 1)] = \mathrm{E}[1 + \delta\Gamma(T_{k,t})] = 1 + \delta\mathrm{E}[\Gamma(T_{k,t})] = 1 + \delta C_{k,t} \ .$$

The future expected cost in slot $t + 1$ with k pending agents can be written in terms of the future expected cost in slot t with k pending agents,

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

∎

A proof of Theorem 7.1 can be found in Appendix C

8 Discussion and Open Problems

Natural directions for future research include the following:

1. Give a protocol for latency costs with expected linear cost (not only with high probability). Alternately, prove that no such protocol exists.
2. Strategic behavior when the number of agents pending is unknown.
3. Allow packets to be inserted over time.
4. Other congestion functions, not only all or nothing. E.g., the outcome of a collision is probabilistic and results in no packet being transmitted with probability p_0, and one random packet being successful with probability p_1. Another example would be where the probability of noise causing a packet to be dropped grows with the number of conflicting transmissions.
5. Consider more general networks, not only a single link.

Perhaps the most interesting direction for future research is to consider generalizations of Wardrop equilibria. Previous work on strategic behavior for multi-commodity 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 work above suggests another parameter for study, the flow duration, *i.e.*, agent i requires flow from s_i to t_i for a duration of d_i. E.g., the flow rate is the bandwidth required to transmit MPEG, the flow duration is the length of the movie. In this setting, it may be advisable to behave politely and allows others to transmit so as to get them out of the way. E.g., it may be advisable to leave home at 10:00 AM so as to avoid rush hour traffic.

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] Joan Feigenbaum, Christos H. Papadimitriou, Rahul Sami, and Scott Shenker. A BGP-based mechanism for lowest-cost routing. In *PODC*, pages 173–182, 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 Osborne and Ariel Rubinstein. *A Course in Game Theory*. Mit., New York, NY, USA, 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*. MIT press, 2005.

A Proof of Theorem 5.1

B Proof of Theorem 5.1

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

The proof goes by induction on k. We show how to prove $\ell_k \leq q_k \leq u_k$ for general lower and upper bounds ℓ_k, u_k.

To show that $\ell_k \leq q_k \leq u_k$, for all k, it suffices to show that

1. If $\ell_k \leq q_k \leq u_k$ then $\frac{1}{(1-\ell_k)^{k-1}} \leq F_k \leq \frac{1}{(1-u_k)^{k-1}}$.
2. If $\frac{1}{(1-\ell_{k-1})^{k-2}} \leq F_{k-1} \leq \frac{1}{(1-u_{k-1})^{k-2}}$ then $\ell_k \leq q_k \leq u_k$.

Proposition B.1. *If* $q_k \geq \ell_k$ *then* $F_k \geq \frac{1}{(1-\ell_k)^{k-1}}$. *If* $q_k \leq u_k$ *then* $F_k \leq \frac{1}{(1-u_k)^{k-1}}$.

Proof. The claim follows directly from equation (5.1), since F_k is monotonically increasing in q_k. ∎

To show the second condition of the induction we 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

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

Consider the rational functions

$$\begin{aligned} \underline{R}(x) &= L(x) - \frac{1}{(1-\ell_{k-1})^{k-2}}, \\ \overline{R}(x) &= L(x) - \frac{1}{(1-u_{k-1})^{k-2}}. \end{aligned}$$

Both $\underline{R}(x)$ and $\overline{R}(x)$ are monotonically increasing in the range $[0, 1]$. Therefore, if $\underline{R}(\ell_k) < 0$ then

$$\underline{R}(x) < 0, x \in [0, 1] \implies x < \ell_k.$$

And similarly, if $\overline{R}(u_k) > 0$ then

$$\overline{R}(x) > 0, x \in [0, 1] \implies x > u_k.$$

We set $\ell_k = \frac{1}{\sqrt{k}+1}$ and $u_k = \frac{2}{\sqrt{k}}$.

Proposition B.2. $\underline{R}(\ell_k) < 0$

Proof. We show that $R(\frac{1}{\sqrt{k-1}+1}) < 0$.

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

From the monotonicity of $\underline{R}(\cdot)$ we get that $R(\frac{1}{\sqrt{k}+1}) \leq R(\frac{1}{\sqrt{k-1}+1}) < 0$. ∎

Proposition B.3. $\overline{R}(u_k) > 0$

Proof. We need to show that

$$\left(\frac{1}{1-u_k}\right)^{k-1}\left(1-\frac{1-u_k}{(k-1)u_k}\right)-\left(\frac{1}{1-u_{k-1}}\right)^{k-2} > 0,$$

or otherwise written

$$\left(\frac{1-u_k}{1-u_{k-1}}\right)^{k-1}(1-u_{k-1})\left(1-\frac{1-u_k}{(k-1)u_k}\right) < 1$$

We present an upper bound of $u_k = \frac{c}{\sqrt{k}}$, where c is a constant that will be chosen later. Hence, we need to show that

$$\left(\frac{1-\frac{c}{\sqrt{k}}}{1-\frac{c}{\sqrt{k-1}}}\right)^{k-1}\left(1-\frac{c}{\sqrt{k-1}}\right)\left(\frac{ck-c}{ck-\sqrt{k}}\right) < 1 \tag{B.1}$$

It can be verified that for every $c > 1$ and large enough k,

$$\left(1-\frac{c}{\sqrt{k-1}}\right)\left(\frac{ck-c}{ck-\sqrt{k}}\right) < 1-\frac{c-\frac{1}{c}}{\sqrt{k}}.$$

The left term in Inequality B.1 can be re-written as

$$\left(\frac{1-\frac{c}{\sqrt{k}}}{1-\frac{c}{\sqrt{k-1}}}\right)^{k-1} = \left(1+\frac{(\sqrt{k}-\sqrt{k-1})c}{\sqrt{k}(\sqrt{k-1}-c)}\right)^{k-1} < \left(1+\frac{c(\frac{1}{2}+\frac{1}{k})}{k(\sqrt{k-1}-c)}\right)^{k-1} \tag{B.2}$$

where the last inequality follows from the bound $\sqrt{k}-\sqrt{k-1} < \frac{1}{2\sqrt{k}}+\frac{1}{k\sqrt{k}}$.

For every c such that $c + 1/c > c/2$ *i.e.*, $c > \sqrt{2}$ and large enough k, it holds that

$$1+\frac{c(\frac{1}{2}+\frac{1}{k})}{k(\sqrt{k-1}-c)} < 1+\frac{c-\frac{1}{c}}{k\sqrt{k}}.$$

As $(1+\frac{1}{y})^y < e$ for all $y \geq 0$, it follows that

$$\left(1+\frac{c-\frac{1}{c}}{k\sqrt{k}}\right)^{k-1} < e^{\frac{(k-1)(c-\frac{1}{c})}{k\sqrt{k}}} < e^{\frac{c-\frac{1}{c}}{\sqrt{k}}},$$

and as $(1+\frac{1}{y}) > e^{\frac{1}{y+1}}$, it follows that

$$e^{\frac{c-\frac{1}{c}}{\sqrt{k}}} < 1+\frac{c-\frac{1}{c}}{\sqrt{k}-(c-\frac{1}{c})} = 1/\left(1-\frac{c-\frac{1}{c}}{\sqrt{k}}\right) \tag{B.3}$$

It follows that for $c = \sqrt{2}+\epsilon$ Inequality B.1 holds every $k \geq M$ for large enough M. This completes the inductive step. To complete the proof we need to provide an induction base.

C Proof of Theorem 7.1

Proof of Theorem 7.1, case $\delta < 1$: First, let us bound $F_{k,t}$, the expected additional cost at state (k,t), from above.

$$F_{k,t} < \sum_{t=t+1}^{\infty} \delta^t = \frac{\delta^{t+1}}{1-\delta}\,. \tag{C.1}$$

Now, let us bound $F_{k,t}$ from below. Consider a miraculous scenario, where there is a successful broadcast on each of the next k slots. Let OPT denote the average future cost on this scenario. Clearly, the true $F_{k,t}$, cannot be less than OPT.

$$
\begin{aligned}
F_{k,t} \geq \mathrm{OPT} &= \frac{1}{k}\sum_{i=0}^{k-1}\Gamma(i+(t+1))-\Gamma(t) \\
&= \frac{1}{k}\sum_{i=0}^{k-1}\sum_{j=t+1}^{i+(t+1)}\delta^j \\
&= \frac{1}{k}\sum_{i=0}^{k-1}\delta^{t+1}\sum_{j=0}^{i}\delta^j \\
&= \frac{\delta^{t+1}}{k}\sum_{i=0}^{k-1}\frac{1-\delta^{i+1}}{1-\delta} \\
&= \frac{\delta^{t+1}}{k}\cdot\frac{1}{1-\delta}\sum_{i=0}^{k-1}1-\delta^{i+1} \\
&= \frac{\delta^{t+1}}{1-\delta}-\frac{\delta^{t+2}}{k(1-\delta)}\sum_{i=0}^{k-1}\delta^i \\
&= \frac{\delta^{t+1}}{1-\delta}-\frac{\delta^{t+2}}{k(1-\delta)}\cdot\frac{1-\delta^k}{1-\delta} \\
&= \frac{\delta^{t+1}}{1-\delta}\left(1-\frac{1}{k}\delta\left(\frac{1-\delta^k}{1-\delta}\right)\right)
\end{aligned}
$$

The last inequality combined with Inequality C.1, implies that, $\frac{F_{k-1,t}}{F_{k,t}} > 1-\frac{1}{k-1}\delta(\frac{1-\delta^{k-1}}{1-\delta})$. Using Lemma 4.2, we get that

$$
q_k = \frac{1}{k-(k-1)\frac{F_{k-1,t}}{F_{k,t}}} > \frac{1}{k-(k-1)(1-\frac{1}{k-1}\delta(\frac{1-\delta^{k-1}}{1-\delta}))} > \frac{1}{1+\frac{\delta}{1-\delta}} = 1-\delta\ .
$$

∎

Proof of Theorem 7.1, case $1<\delta<e/(e-1)$: Recall the definitions of $\alpha_{k,t},\beta_{k,t}$ (Equations 4.1and 4.2). The time index t is redundant in the description of $\alpha_{k,t},\beta_{k,t}$, since time independent strategies are considered. Hence we will write α_k,β_k instead. The expected cost

$$
C_{k,t} = \alpha_k\cdot\Gamma(t+1)+\beta_k C_{k-1,t+1}+(1-\alpha_k-\beta_k)C_{k,t+1}. \tag{C.2}
$$

Using Lemma 7.2 we can remove the dependence on time in Equation C.2,

$$
C_{k,t} = \alpha_k\cdot(1+\delta\Gamma(t))+\beta_k(1+\delta C_{k-1,t})+(1-\alpha_k-\beta_k)(1+\delta C_{k,t}).
$$

Rearranging we get

$$
(C_{k,t}-\Gamma(t))(1-\delta(1-(\alpha_k+\beta_k)))-\delta^{t+1} = \delta\beta_k(C_{k-1,t}-\Gamma(t))\ ,
$$

and since $\delta^t > 0$

$$(C_{k,t} - \Gamma(t))(1 - \delta(1 - (\alpha_k + \beta_k))) > \delta\beta_k(C_{k-1,t} - \Gamma(t)) \ .$$

We replace expected cost with future expected cost and get that

$$\frac{F_{k-1,t}}{F_{k,t}} = \frac{C_{k-1,t} - \Gamma(t)}{C_{k,t} - \Gamma(t)} < \frac{(1 - \delta(1 - (\alpha_k + \beta_k)))}{\delta\beta_k} \ .$$

Replacing α_k, β_k as a function of q_k, we get an upper bound on $\frac{F_{k-1,t}}{F_{k,t}}$ as a function of q_k and k.

$$\frac{F_{k-1,t}}{F_{k,t}} < \frac{1 - \delta(1 - kq_k(1 - q_k)^{k-1})}{\delta(k-1)q_k(1 - q_k)^{k-1}} \ . \tag{C.3}$$

Let $g(q) = q(1-q)^{k-1}$, and let $f(x) = \frac{1-\delta(1-kx)}{\delta(k-1)x}$. The right hand side of Inequality C.3 equals $f(g(q_k))$. We show now an upper bound on $f(g(q))$ in the range $q \in [0,1]$. Note that $\mathrm{argmax}_{q\in[0,1]} g(q) = 1/k$. Therefore, $g(q) \leq \frac{1}{k-1}(1 - \frac{1}{k})^k < \frac{1}{e(k-1)}$, for every $0 \leq q \leq 1$. For $x < \frac{1}{e(k-1)}$,

$$\begin{aligned} f(x) &= \frac{1 - \delta(1 - kx)}{\delta(k-1)x} \\ &= \frac{k}{k-1} - \frac{\delta - 1}{\delta(k-1)x} \\ &< \frac{k}{k-1} - \frac{\delta - 1}{\delta(k-1)\frac{1}{e(k-1)}} \\ &< \frac{k}{k-1} - \frac{e(\delta - 1)}{\delta} \end{aligned}$$

Using Lemma 4.2, we get that

$$q_k < \frac{1}{k - (k-1)(\frac{k}{k-1} - \frac{e(\delta-1)}{\delta})} = \frac{1}{k-1} \cdot \frac{\delta}{e(\delta - 1)}$$

∎