

Contention Issues in Congestion Games^{*}

Elias Koutsoupias and Katia Papakonstantinou

University of Athens, Athens, Greece
{elias,katia}@di.uoa.gr

Abstract. We study time-dependent strategies for playing congestion games. The players can time their participation in the game with the hope that fewer players will compete for the same resources. We study two models: the boat model, in which the latency of a player is influenced only by the players that start at the same time, and the conveyor belt model in which the latency of a player is affected by the players that share the system, even if they started earlier or later; unlike standard congestion games, in these games the order of the edges in the paths affect the latency of the players. We characterize the symmetric Nash equilibria of the games with affine latencies of networks of parallel links in the boat model and we bound their price of anarchy and stability. For the conveyor belt model, we characterize the symmetric Nash equilibria of two players on parallel links. We also show that the games of the boat model are themselves congestion games. The same is true for the games of two players for the conveyor belt model; however, for this model the games of three or more players are not in general congestion games and may not have pure equilibria.

Keywords: Algorithmic game theory, price of anarchy, congestion games, contention.

1 Introduction

In the last dozen years, the concepts of the price of anarchy (PoA) and stability (PoS) have been successfully applied to many classes of games, most notably to congestion games and its relatives [17,24,21]. In congestion games, the players compete for a set of resources, such as facilities or links; the cost of each player depends on the number of players using the same resources; the assumption is that each resource can be shared among the players, but with a cost. Another interesting class of games are the contention games [12] in which the players again compete for resources, but the resources cannot be shared. If more than one players attempt to share a resource at the same time, the resource becomes unavailable and the players have to try again later. There are however interesting games that lie between the two extreme cases of the congestion and contention games. For example, the game that users play for dealing with congestion on a network seems to lie in between—the TCP congestion control policy is a strategy

^{*} Partially supported by the ESF-NSRF research program Thales (AGT).

of this game. Timing is part of the strategy of the players (as in contention games) and the latency of a path depends on how many players use its edges (as in congestion games).

In this work, we attempt to abstract away the essential features of these games, to model them, and to study their properties, their Nash equilibria, and their price of anarchy and stability. The games that we consider are essentially congestion games with the addition of time dimension. The difference with congestion games is that players now don't simply select which path to use, but they also decide *when* to initiate the transmission.

Consider a link or facility e of a congestion game with latency function ℓ_e . In the congestion game the latency that a player experiences on the link is $\ell_e(k)$, where k is the number of players that use the link. In our model however, in which the players can also decide when to start, the latency needs to be redefined. We define and study two latency models for the links:

The boat model: in which only the group of players that start together affect the latency of the group: imagine that one boat departs from the source of the link at every time step; all players that decide to start at time t enter the boat which takes them to their destination; the speed of the boat depends only on the number of players in the boat and it is independent of the players on the other boats.

The conveyor belt model: in which the latency of a player depends on the number of other players using the link at the same time regardless if they started earlier or later. Specifically, the link is like a conveyor belt from the source to the destination; the speed of the belt at every time depends on the number of people on it. An interesting variant of this model is when the player is affected only by the players that have been already in the link but not by the players that follow; we don't study this model in this work.

Notice that in the boat model, the order in which the players finish a link may differ from the order in which they start. This, for example, can happen when a player starts later but with a smaller group of people. This cannot happen in the conveyor belt model.

In this work, we consider

- non-adaptive strategies, in which the players decide on their strategy in advance. Their pure strategies consist of a path and a starting time.
- symmetric strategies.

Intuitively, in the boat model, the aim of the players is to select a path with small latency and to avoid other players that start at the same time. In the conveyor belt model the aim is similar but the players try to avoid other players that start *near* the same time.

Related work. Contention resolution in communication networks is a problem that has attracted the interest of diverse communities of Computer Science. Its significance comes from the fact that contention is inherent in many critical

network applications. One of them is the design of multiple access protocols for communication networks, such as Slotted Aloha: According to it, a source transmits a packet through the network, as soon as this packet is available. If a collision takes place, that is, another source attempted to transmit simultaneously, the source waits for some random number of time slots and attempts to retransmit at the beginning of the next slot. The increase of users of the network incurs a large number of collisions and subsequently poor utilization of the system's resources.

During the last four decades many more refined multiple access protocols have been proposed to increase the efficiency of Aloha, the vast majority of which assume that the agents follow the protocol, even if they might prefer not doing so. Recently, slotted Aloha has been studied from a game-theoretic point of view, trying to capture the selfish nature of its users. Part of this work has been done by Altman et al. [2,3]. The authors model slotted Aloha as a game among the transmitters who aim at transmitting their stochastic flow, using the retransmission probability that maximizes their throughput [2] or minimizes their delay [3]. They show that the system possesses symmetric equilibria and that its throughput deteriorates with larger number of players or arrival rate of new packets. Things get better considering a cost for each transmission though. Another slotted Aloha game is studied by MacKenzie and Wicker [18]. Here the agents aim at minimizing the time spent for unsuccessful transmissions before each successful one, while each transmission incurs some cost to the player. Their game possesses a symmetric equilibrium, and some of its instantiations possess equilibria that achieve the maximum possible throughput of Aloha.

Much of the prior game-theoretic work considers transmission protocols that always transmit with the same fixed probability. In [12] and [11] the authors consider more complex protocols (multi-round games), where a player's transmission probability is allowed to be an arbitrary function of his play history and the sequence of feedback he has received, and propose asymptotically optimal protocols. In [12], the authors propose a protocol which is a Nash equilibrium and has constant price of stability, i.e., all agents will successfully transmit within time proportional to their number. This protocol assumes that the cost of any single transmission is zero. In [11] the case of non-zero transmission cost is addressed, and a protocol is proposed where after each time slot, the number of attempted transmissions is returned as feedback to the users.

There is a lot of work on game theoretic issues of packet switching. For example, [15] considers the game in which users select their transmission rate, [1] considers TCP-like games in which the strategies of the players are the parameters of the AIMD (additive increase / multiplicative decrease) algorithm, and [13] considers game-theoretic issues of congestion control. All these works are concerned with the steady or long term version of the problems and they don't consider time-dependent strategies in the spirit of this work.

Routing in networks by selfish agents is another area that has been extensively studied based on the notion of the price of anarchy (PoA) [17] and the price of stability (PoS) [5]. The PoA and the PoS compare the social cost of

the worst-case and best-case equilibrium to the social optimum. Selfish routing is naturally modeled as a congestion game. The class of congestion or potential games [22,20] consists of the games where the cost of each player depends on the resources he uses and the number of players using each resource. The effect of selfishness in infinite congestion games was first studied in [24] and of finite congestion games in [8,6].

The above results concern classical networks or static flows on networks. Perhaps the closest in spirit to our work are the recent attempts to study game-theoretic issues of dynamic flows, or more precisely, of flows over time. In [16], the authors consider selfish selection of routing paths when users have to wait in a FIFO queue before using every edge of their paths; the waiting time is not part of their strategy, but depends on the traffic in front of them. The same model is assumed by [19] who considers the Braess' paradox for flows over time. More results appeared in [7] which gives an efficiently computable Stackelberg strategy for which the competitive equilibrium is not much worse than the optimal, for two natural measures of optimality: total delay of the players and time taken to route a fixed amount of flow to the sink. In a slightly different model, [4] considers game-theoretic issues of discrete-time models in which the latency of each edge depends on its history. All these papers consider non-atomic congestion games. In a different direction which involves atomic games, [14] considers temporal congestion games that are based on coordination mechanisms [10] and congestion games with time-dependent costs.

All these models share with this work the interest in game-theoretic issues of timing in routing, but they differ in an essential ingredient: in our games, timing is the most important part of the players strategy, while in the previous work, time delays exist because of the interaction of the players; in particular, *in all these models the strategy of the players is to select only a path*, while in our games the strategy is essentially the timing. We view our model as a step towards understanding games related to TCP congestion control; this does not seem to be in the research agenda of game-theoretic issues of flows over time.

Short description of results. We first study structural properties of the boat and conveyor belt games. In the next section, we characterize the symmetric Nash equilibria and the optimal symmetric solution of the boat model game for parallel links of affine latency functions and any number of players. From these we get that the price of anarchy and stability is very low $3\sqrt{2}/4 \approx 1.06$. We also study the class of conveyor belt games. These are more complicated games and here we consider only two players and arbitrary latency functions (for two players the class of affine and the class of arbitrary latency functions are identical). The price of anarchy and stability is (for large latencies) again approximately $3\sqrt{2}/4 \approx 1.06$. This is the price of anarchy we computed for the boat model, but the relation is not as straightforward as it may appear: in the boat model we take the limit as the number of players tends to infinity, while in the conveyor model, we take the limit as the latencies tend to infinity.

To our knowledge, these games differ significantly from the classes of congestion games that have been studied before. Also, the techniques developed for

bounding the PoA and the PoS for congestion games do not seem to be applicable in our setting. In particular, the smoothness analysis arguments [8,23,9] do not seem to apply because we consider symmetric equilibria. In fact, the focus and difficulty of our analysis is to characterize the Nash equilibria and not to bound the PoA (or PoS).

The decision to study only symmetric strategies is based on the assumption that these games are played by many players with no coordination among them. We consider this work as a step towards the study of real-life situations such as the TCP congestion control mechanism in which the players are essentially indistinguishable and therefore symmetric.

In all the games that we study, there exists a unique symmetric equilibrium. For this type of equilibria, the definition of the price of anarchy is uncomplicated: We simply take the ratio of the cost of one player over the cost of one player of *the symmetric optimal solution*. Since there is a unique Nash equilibrium, the price of stability is equal to the price of anarchy.

Due to the space limitations, some proofs are omitted but can be found in the full version of the paper. Moreover, the structural properties of the two models are presented in short here, and are described in detail in the full version of this work.

Formally, the games that we study here are the following: Let G be a *network* congestion game with n players and latency functions $\ell_e(k)$ on its link e . We define two new games based on G , the boat model game and the conveyor belt game. The pure strategies of both new games of every player consist of one strategy (path) of the original game and one non-negative time step $t \in \mathbb{Z}_0^+$. Their difference lies in the cost of the pure strategies.

In the boat model, the cost of a player is simply $t + \sum_{e \in P} \ell_e(n_t(e))$, where $n_t(e)$ denotes the set of players that also start at time t and use edge e . In the conveyor belt model the cost is more complicated. It depends on the notion of work: in a time interval $[t, t + \Delta t]$ in which player i uses link e , it completes work $\Delta t / \ell_e(k)$, where k is the number of players using the same link during this time interval. A player finishes a link when it completes total work of 1 for this link; the player then moves to the next link of its path.

The following theorem describes the nature of the time-dependent games.

Theorem 1. *All boat games are congestion games. In contrast, only the 2-player conveyor belt games are congestion games, and for 3 or more players there are games that have no pure equilibria. Furthermore, the order of using the facilities in conveyor belt games is important: a reordering of the edges of a path can result in a different game.*

2 Nash Equilibria of the Boat Model

In this section, we first consider symmetric Nash equilibria of n players for the boat model of parallel links. We also compute the optimal non-selfish solution and estimate the PoA.

Nash equilibria computation. A pure strategy for a player is to select a link e and a time t . A mixed strategy is given by probabilities $p_{e,t}$ with $\sum_{e,t} p_{e,t} = 1$: the player uses link e at time step t with probability $p_{e,t}$. A set of probabilities $p_{e,t}$ is a Nash equilibrium when a player has no incentive to change it to some other values q . To find the Nash equilibria, we first estimate the latency $d_{e,t}$ when the player selects pure strategy (e, t) :

$$d_{e,t} = t + \sum_{k=0}^{n-1} \binom{n-1}{k} p_{e,t}^k (1-p_{e,t})^{n-1-k} \ell_e(k+1). \quad (1)$$

Let $d = \min_{e,t} d_{e,t}$ denote the minimum value. Then the probabilities define a symmetric mixed Nash equilibrium if and only if $p_{e,t} > 0$ implies $d = d_{e,t}$.

To find the Nash equilibria, the first crucial step is to show that the probabilities in every link must be non-increasing in t . This is shown by the following lemma which holds for arbitrary latency functions, not only for affine ones:

Lemma 1. *If for every edge e the latencies $\ell_e(k)$ are non-decreasing in k , then every symmetric Nash equilibrium is a non-increasing sequence of probabilities: $p_{e,t} \geq p_{e,t+1}$.*

(The proof is omitted and can be found in the full version of the paper.)

We define the support of the Nash equilibrium to be the set of strategies that have minimum latency: $S_e = \{t : d_{e,t} = d\}$. Alternatively, we could have defined the support to be the set of strategies with non-zero probability at the Nash equilibrium; the two notions are similar but not identical in some cases. Notice the convention $d_{e,h_e+1} > d = d_{e,h_e}$, in the definition of the support. The last lemma shows that the support S_e of every link e is of the form $\{0, \dots, h_e\}$ for some integer h_e .

We now focus on affine latency functions, $\ell_e(k) = a_e k + b_e$, for which the cost $d_{e,t}$ in (1) takes a simple closed form:

$$d_{e,t} = t + a_e + b_e + (n-1) a_e p_{e,t}, \quad (2)$$

which shows that the probabilities of the Nash equilibria are of the form:

$$p_{e,t} = \begin{cases} \frac{d-a_e-b_e-t}{(n-1)a_e} & \text{for } t \leq h_e \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Observe that at every Nash equilibrium $p_{e,t}$, the non-zero probabilities decrease linearly with t . These probabilities are determined by the cost d of each player and the integers h_e (one for each link). In fact, the parameters h_e are very tightly related with the cost d of each player:

Theorem 2. *There is a unique symmetric Nash equilibrium with support $S_e = \{t : 0 \leq t \leq h_e = \lfloor d - a_e - b_e \rfloor\}$, where d is the expected cost of every player; its probabilities are given by*

$$p_{e,t} = \begin{cases} \frac{d-a_e-b_e-t}{(n-1)a_e} & \text{for } t \leq d - a_e - b_e \\ 0 & \text{otherwise} \end{cases}$$

The expected cost $L_{NE} = d$ of every player is the unique solution of the equation

$$\sum_e \frac{(\lfloor d - a_e - b_e \rfloor + 1)(2(d - a_e - b_e) - \lfloor d - a_e - b_e \rfloor)}{2(n-1)a_e} = 1. \quad (4)$$

Its value is approximately

$$d \approx \frac{\sum_e \frac{a_e + b_e}{2(n-1)a_e} + \sqrt{\left(\sum_e \frac{a_e + b_e}{2(n-1)a_e}\right)^2 + \left(\sum_e \frac{1}{2(n-1)a_e}\right) \left(1 - \sum_e \frac{(a_e + b_e)^2}{2(n-1)a_e}\right)}}{\sum_e \frac{1}{2(n-1)a_e}}, \quad (5)$$

and as n tends to infinity this tends to $\sqrt{\frac{2n}{\sum_e a_e^{-1}}}$.

(The proof is omitted and can be found in the full version of the paper.)

The optimal setting. Let us now consider the optimal symmetric protocol. With similar reasoning, the expected latency of a player is

$$L_{OPT} = \sum_e \sum_{t=0}^{\infty} p_{e,t} \left(t + \sum_{k=0}^{n-1} \binom{n-1}{k} p_{e,t}^k (1 - p_{e,t})^{n-1-k} \ell_e(k+1) \right) = \sum_e \sum_{t=0}^{\infty} p_{e,t} d_{e,t}$$

We seek the probabilities $p_{e,t}$ with $\sum_{e,t} p_{e,t} = 1$ which minimize the above expression. We again focus on affine latencies. With $\ell_e(k) = ak + b$, the above expression has the following compact form $L_{OPT} = \sum_e \sum_{t=0}^{\infty} p_{e,t} (t + a_e + b_e + (n-1)a_e p_{e,t})$. We minimize this subject to $\sum_{e,t} p_{e,t} = 1$. Using a Lagrange multiplier and taking derivatives, we get that the minimum occurs when the probabilities have the form $p_{e,t} = (\lambda - a_e - b_e - t)/(2(n-1)a_e)$, for some constant λ , and $p_{e,t} = 0$ when $\lambda - a_e - b_e - t \leq 0$. This means that they decrease linearly with t until $c_e = \lambda - a_e - b_e$, when they become 0 and they remain 0 from that point on. Thus, the form of the optimal probabilities resembles the form of the Nash equilibrium probabilities; the only difference is that the optimal probabilities drop slower to 0 (the factors are $2(n-1)a_e$ and $(n-1)a_e$ respectively). A similar bicriteria relation between the Nash equilibria and the optimal solution has been observed in simple congestion games before [24]. Taking into account the constant term also we get,

Lemma 2. *The set of probabilities of the optimal solution for latencies $\ell_e(k) = a_e k + b_e$ is a Nash equilibrium for latencies $\ell_e(k) = 2a_e k + (b_e - a_e)$.*

Therefore the probabilities of the optimal solution are:

$$p_{e,t} = \begin{cases} \frac{\lambda - a_e - b_e - t}{2(n-1)a_e} & t \leq h_e^* \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where $h_e^* = \lfloor \lambda - a_e - b_e \rfloor$, and the value of λ is the unique solution of the equation $\sum_{e,t} p_{e,t} = 1$. Thus, λ is determined by an equation similar to (4)

(they essentially differ only in the denominator):

$$\sum_e \frac{([\lambda - a_e - b_e] + 1)(2(\lambda - a_e - b_e) - [\lambda - a_e - b_e])}{4(n-1)a_e} = 1. \quad (7)$$

From the probabilities we can compute L_{OPT} . Observe that the optimal case differs from the Nash equilibrium case of the previous subsection in that the parameters λ and L_{OPT} are distinct (while in the Nash equilibrium case they are identical—equal to d).

As in the case of the Nash equilibrium, it is useful to define $\eta_e^* = \lambda - a_e - b_e$. We can then compute the optimal latency: $L_{OPT} =$

$$\sum_e \sum_{t=0}^{\infty} p_{e,t} (t + a_e + b_e + (n-1)a_e p_{e,t}) = \sum_e \frac{(h_e^* + 1)(6\eta_e^*(\eta_e^* + 2a_e) - h_e^*(2h_e^* + 6a + 1))}{24(n-1)a_e}.$$

To get an approximate estimate as n tends to infinity, we observe that λ is approximately given by $\sum_e \frac{\lambda^2}{4(n-1)a_e} \approx 1$ which implies $\lambda \approx 2\sqrt{\frac{n}{\sum_e a_e^{-1}}}$. From this, we can find an approximate value for L_{OPT} :

$$L_{OPT} \approx \frac{\eta_e^{*3}}{6(n-1)a_e} \approx \sum_e \frac{\lambda^3}{6(n-1)a_e} = \frac{4}{3} \sqrt{\frac{n}{\sum_e a_e^{-1}}}$$

The price of anarchy. Comparing the value of L_{OPT} to the cost d of the Nash equilibrium, we see that the PoA and the PoS of the boat model on parallel links with affine latency functions tends to $\frac{3\sqrt{2}}{4} \approx 1.06$, as the number of players n tends to infinity (while the parameters of the network remain fixed).

Theorem 3. *For every fixed set of parallel links with positive a_e and b_e , the PoA (and PoS) tends to $3\sqrt{2}/4 \approx 1.06$, as the number of players n tends to infinity.*

However, for fixed number of players and because of the integrality of h and h^* , the situation is more complicated. Figure 1 shows the PoA for typical values of a_e and n , for one link. The situation is captured by the following theorem:

Theorem 4. *For one link and fixed number of players n , the PoA is maximized when $a_e = 1/(n-1)$ and $b_e = 0$. For these values the NE is pure ($p_{e,0} = 1$), but the optimal symmetric solution is given by the probabilities $p_{e,0} = 3/4$ and $p_{e,1} = 1/4$. For these strategies we get $L_{NE} = d = n/(n-1)$, $L_{OPT} = (7n+1)/(8(n-1))$, and $PoA = 8n/(7n+1)$.*

To compare the cost L_{NE} and L_{OPT} we first investigate the solutions of the equations (4) and (7) as functions of a_e ; since we care about the worst-case PoA, we can safely assume that $b_e = 0$ because $b_e \geq 0$ is added to both the numerator and the denominator of the PoA.

For every nonnegative integer k , let us define $A_k = \frac{k(k+1)}{2(n-1)}$ which are the values of a_e where the value $d - a_e$ becomes integral (equal to k). The following lemma gives the solution of (4) for the intervals $[A_k, A_{k+1})$ where the integral part of $d - a_e$ is constant. It also extends it to the optimal cost.

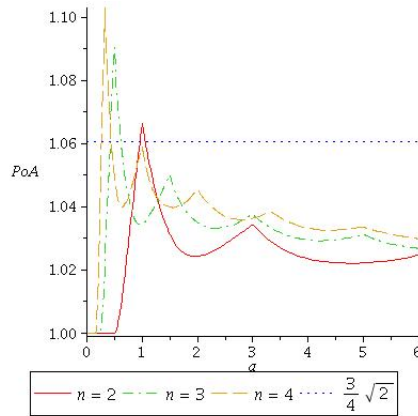


Fig. 1. PoA of the single-link boat games

Lemma 3. For $a_e \in [A_k, A_{k+1})$, $L_{NE} = \frac{n+k}{k+1}a_e + \frac{k}{2}$.
 For $a_e \in [A_k/2, A_{k+1}/2)$, $L_{OPT} = \frac{n+k}{k+1}a_e + \frac{k}{2} - \frac{k(k+1)(k+2)}{48(n-1)a_e}$.

Proof. We first show that for $a_e \in [A_k, A_{k+1})$ the value of d given by equation (4) satisfies $\lfloor d - a_e \rfloor = k$. It suffices to show that the value $x = d - a_e - k$ is in $[0, 1)$.

3 Nash Equilibria of the Conveyor Belt Model

We now turn our attention to the conveyor belt model, which is more complicated than the boat model. In the conveyor belt model each link is like a conveyor belt whose speed depends on the number of players on it. *We only consider the case of 2 players in this section.* The cost $c_e(t, t')$ of a player for pure strategies (e, t) when the other player starts using link e at time step t' is computed using $f_i = t_i + \ell_e(1) + \max\left(0, (\ell_e(2) - \ell_e(1)) \left(1 - \frac{|t_2 - t_1|}{\ell_e(1)}\right)\right)$ where t_i, f_i are the start and finish times of player i respectively (see the full version for details).

To simplify the discussion, we assume that $\ell_e(1)$ is an integer; this does not seem to really change the nature of equilibria, except perhaps when $\ell_e(1) < 1$ which does not seem a very interesting case.

Nash equilibria computation. Consider a symmetric Nash equilibrium with probabilities $p_{e,t}$, the same for every player. It is a Nash equilibrium when a player has no incentive to change his probabilities to different values. To find the Nash equilibria, we first compute the expected cost $d_{e,t}$ of a player when he plays pure strategy (e, t) :

$$d_{e,t} = \sum_{t'} c_e(t, t') = t + \ell_e(1) + (\ell_e(2) - \ell_e(1)) \sum_{r=-\ell_e(1)}^{\ell_e(1)} \left(1 - \frac{|r|}{\ell_e(1)}\right) p_{e,t+r} \quad (8)$$

The probabilities define a symmetric mixed Nash equilibrium when probability $p_{e,t} > 0$ implies $d_{e,t} = d = \min_{e,t} d_{e,t}$.

We are interested in symmetric Nash equilibria, that is equilibria that occur when all players use the same strategies. Let's first establish a very intuitive fact:

Claim. If at the Nash equilibrium, positive probability is allocated to edge e , then $p_{e,0} > 0$.

(The proof is omitted and can be found in the full version of the paper.)

The next lemma shows that the support $S_e = \{t : d_{e,t} = d\}$ of every mixed Nash equilibrium is of the form $\{0, \dots, \hat{h}_e\}$ for some \hat{h}_e .

Lemma 4. *If for some t there exists $s \geq t$ with $p_{e,t} > 0$, then t is in the support S_e , i.e. $d_{e,t} = d$.*

(The proof is omitted and can be found in the full version of the paper.)

The previous lemma establishes that the support S_e starts at 0 and is contiguous. With this, we can now determine the exact structure of Nash equilibria.

Theorem 5. *The Nash equilibria of the conveyor belt game of two players in parallel links have probabilities*

$$p_{e,t} = \begin{cases} \frac{d - \ell_e(1) - t}{\ell_e(2) - \ell_e(1)} & t \leq d - \ell_e(1) \text{ and } \frac{t}{\ell_e(1)} \in \mathbb{Z}^+ \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

where d is the expected cost of each player and it is the unique solution of the equation

$$\sum_e \frac{(\lfloor \eta_e \rfloor + 1)(2\eta_e - \lfloor \eta_e \rfloor)}{2^{\frac{\ell_e(2) - \ell_e(1)}{\ell_e(1)}}} = 1, \quad (10)$$

where $\eta_e = d/\ell_e(1) - 1$.

Proof. Consider some $0 < t < \hat{h}_e$. Then from the definition of $d_{e,t}$ we can compute $d_{e,t+1} - 2d_{e,t} + d_{e,t-1} = \frac{\ell_e(2) - \ell_e(1)}{\ell_e(1)}(p_{e,t-\ell_e(1)} - 2p_{e,t} + p_{e,t+\ell_e(1)})$. Since for $t \in \{1, \dots, \hat{h}_e - 1\}$, all $t-1$, t and $t+1$ are in the support S_e , we have that $d_{e,t-1} = d_{e,t} = d_{e,t+1}$. In turn, this gives that the right-hand side is 0 and we get that $p_{e,t+\ell_e(1)} - p_{e,t} = p_{e,t} - p_{e,t-\ell_e(1)}$; this shows that if we consider times that differ by $\ell_e(1)$, the probabilities drop linearly and more specifically that for integers k, x : $p_{e,k\ell_e(1)+x} - p_{e,x} = k(p_{e,x+\ell_e(1)} - p_{e,x})$.

This linearity allows us to conclude that $p_{e,t} = 0$ for every t which is not a multiple of $\ell_e(1)$. To see this consider some $x \in \{1, \dots, \ell_e(1) - 1\}$ and the sequence $p_{e,x-\ell_e(1)}, p_{e,x}, p_{e,x+\ell_e(1)}, \dots, p_{e,x+k\ell_e(1)}$. This sequence is linear and starts with a 0 (since $x - \ell_e(1) < 0$) and ends again in 0 (if we take k such that $\hat{h}_e < x + k\ell_e(1) \leq \hat{h}_e + \ell_e(1)$).

The above reasoning does not apply to the value $x = 0$, because $p_{e,t+\ell_e(1)} - p_{e,t} = p_{e,t} - p_{e,t-\ell_e(1)}$ only for $t \in \{1, \dots, \hat{h}_e - 1\}$. To summarize, the NE with support $\{0, \dots, \hat{h}_e\}$ have non-zero probabilities only on the multiples of $\ell_e(1)$.

This means that either the players start together, or they do not overlap, which is *exactly the property of the boat model*. It follows, that for one link, the Nash equilibrium is identical to the Nash equilibrium of the boat game with time step expanded to $\ell_e(1)$. For more than one link, the time steps in each link are different, because $\ell_e(1)$ are different. Nevertheless the analysis of the boat model carries over to the conveyor belt model.

The proof now is essentially the same with the boat model, but with the extra restriction that the time steps are not the same in all links. Since the probabilities are non-zero only at integral multiples of $\ell_e(1)$, the latency becomes $d_{e,t} = t + \ell_e(1) + (\ell_e(2) - \ell_e(1))p_{e,t}$ when t is an integral multiple of $\ell_e(1)$. It follows that the probabilities are as in (9). The cost d is determined by the equation $\sum_{e,t} p_{e,t} = 1$. Using the expressions for the probabilities, this equation is equivalent to (10). This is identical to the equation for d for the boat model and the argument about the uniqueness of the solution carries over.

The Optimal setting. Let's now consider the optimal symmetric protocol. With reasoning similar to that in the boat model and omitting the details (which can be found in the full version of the paper), we get that the minimum occurs when

$$\lambda = t + \ell_e(1) + 2(\ell_e(2) - \ell_e(1)) \sum_{r=-\ell_e(1)}^{\ell_e(1)} \left(1 - \frac{|r|}{\ell_e(1)}\right) p_{e,t+r}, \quad (11)$$

for some λ . The factor 2 in the last term comes from the convolution in the L_{OPT} expression. We notice again the bicriteria property.

Lemma 5. *The probabilities of the optimal solution for two players in the conveyor belt model of parallel links with latencies $\ell_e(k)$ is a Nash equilibrium for latencies $\ell'_e(k) = 2\ell_e(k) - \ell_e(1)$.*

Proof. By comparing equations (8) and (11) that determine the Nash equilibria and the optimal solution, we see that the latencies must satisfy $\ell'_e(1) = \ell_e(1)$ and $\ell'_e(2) - \ell'_e(1) = 2(\ell_e(2) - \ell_e(1))$, which can be expressed as in the lemma.

Since the conveyor belt Nash equilibrium and optimal solution are very similar to the ones of the boat model, the analysis of the price of anarchy is similar, and their expressions can be approximated well as the latencies $\ell_e(k)$ tend to infinity. For one link, the cost d of the Nash equilibrium is approximately $\sqrt{2\ell_e(1)(\ell_e(2) - \ell_e(1))}$ while the optimal cost is $\frac{4}{3}\sqrt{2\ell_e(1)(\ell_e(2) - \ell_e(1))}$, which shows that the price of anarchy tends to $3\sqrt{24} \approx 1.06$, again. Since this is not sufficiently different than the boat model, we omit the details.

References

1. A. Akella, S. Seshan, R. Karp, S. Shenker, and C. Papadimitriou. Selfish behavior and stability of the internet: a game-theoretic analysis of TCP. In *Proceedings of the 2002 SIGCOMM conference*, pages 117–130. ACM, 2002.

2. Eitan Altman, Rachid El Azouzi, and Tania Jiménez. Slotted aloha as a game with partial information. *Comput. Netw.*, 45:701–713, August 2004.
3. Eitan Altman, Dhiman Barman, Rachid El Azouzi, and Tania Jiménez. A game theoretic approach for delay minimization in slotted ALOHA. In *IEEE International Conference on Communications*, 2004.
4. Elliot Anshelevich and Satish Baruch. Equilibria in dynamic selfish routing. In *Proceedings of SAGT '09*, pages 171–182, Berlin, Heidelberg, 2009. Springer-Verlag.
5. Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of FOCS '04*, pages 295–304, 2004.
6. Baruch Awerbuch, Yossi Azar, and Amir Epstein. Large the price of routing un-splittable flow. In *Proceedings of STOC '05*, pages 57–66, 2005.
7. Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A stackelberg strategy for routing flow over time. In *Proceedings of SODA '11*, pages 192–201. SIAM, 2011.
8. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of STOC '05*, pages 67–73. ACM, 2005.
9. G. Christodoulou, E. Koutsoupias, and P. Spirakis. On the performance of approximate equilibria in congestion games. In *European Symposium on Algorithms*, pages 251–262, Copenhagen, Denmark, 7-9 September 2009. Also in arXiv:cs/0804.3160.
10. George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
11. George Christodoulou, Katrina Ligett, and Evangelia Pyrga. Contention resolution under selfishness. In *Proceedings of ICALP 2010*, pages 430–441. Springer-Verlag.
12. Amos Fiat, Yishay Mansour, and Uri Nadav. Efficient contention resolution protocols for selfish agents. In *Proceedings of SODA '07*, pages 179–188, 2007.
13. R. Garg, A. Kamra, and V. Khurana. A game-theoretic approach towards congestion control in communication networks. *ACM SIGCOMM Computer Communication Review*, 32(3):47–61, 2002.
14. Martin Hoefer, Vahab S. Mirrokni, Heiko Röglin, and Shang-Hua Teng. Competitive routing over time. In *Proceedings of WINE 2009*, pages 18–29. Springer-Verlag.
15. A. Kesselman, S. Leonardi, and V. Bonifaci. Game-theoretic analysis of internet switching with selfish users. *Internet and Network Economics*, pages 236–245, 2005.
16. Ronald Koch and Martin Skutella. Nash equilibria and the price of anarchy for flows over time. In *Proceedings of SAGT '09*, pages 323–334. Springer-Verlag, 2009.
17. Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of STACS '99*, pages 404–413, Berlin, Heidelberg, 1999. Springer-Verlag.
18. Allen B. MacKenzie and Stephen B. Wicker. Stability of multipacket slotted aloha with selfish users and perfect information. In *Proceedings of IEEE INFOCOM*, pages 1583–1590, 2003.
19. Martin Macko, Kate Larson, and L'uboš Steskal. Braess's paradox for flows over time. In *Proceedings of SAGT '10*, pages 262–275. Springer-Verlag, 2010.
20. D. Monderer and L.S. Shapley. Potential games. *Games and economic behavior*, 14:124–143, 1996.
21. N. Nisan, T. Roughgarden, É. Tardos, and V.V. Vazirani. *Algorithmic game theory*. Cambridge Univ Pr, 2007.
22. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
23. Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of STOC '09*, pages 513–522, New York, NY, USA, 2009. ACM.
24. Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, March 2002.