

Games with Congestion-Averse Utilities

Andrew Bye¹, Maria Polukarov², and Nicholas R. Jennings²

¹ Hewlett-Packard Laboratories, Bristol, UK
`andrew.bye@hp.com`

² School of Electronics and Computer Science,
University of Southampton, UK
`mp3,nrj@ecs.soton.ac.uk`

Abstract. Congestion games—in which players strategically choose from a set of “resources” and derive utilities that depend on the congestion on each resource—are important in a wide range of applications. However, to date, such games have been constrained to use utility functions that are linear sums with respect to resources. To remove this restriction, this paper provides a significant generalisation to the case where a player’s payoff can be given by *any* real-valued function over the set of possible congestion vectors. Under reasonable assumptions on the structure of player strategy spaces, we constructively prove the existence of a pure strategy equilibrium for the very wide class of these generalised games in which player utility functions are *congestion-averse*—i.e., monotonic, submodular and independent of irrelevant alternatives. Although, as we show, these games do not admit a generalised ordinal potential function (and hence—the finite improvement property), any such game does possess a Nash equilibrium in pure strategies. A polynomial time algorithm for computing such an equilibrium is presented.

1 Introduction

Models of congestion have recently become a major issue of study in algorithmic game theory, as they arise from many real-life situations (examples include network routing, resource and task allocation, competition of firms for production processes [8, 16, 17]) and yet possess plausible theoretical properties. To date, much of research deals with the model of *congestion games* introduced by Rosenthal [16]. Here, players share a finite set of resources, and a strategy for a player is to choose a subset of the resources. Each resource is associated with a resource utility function, which determines the utility of each of its users as a function of the number of players that have selected the resource. Given a strategy profile, a combination of the players’ chosen strategies, the payoff for a player will be simply the sum of utilities from his utilised resources.

Now, the key result of Rosenthal is that congestion games always possess pure strategy Nash equilibria. This is important because pure strategy equilibria have some indisputable advantages over mixed strategy equilibria: they are more intuitive, especially in the context of one-shot games, they are generally easier to compute than mixed equilibria, and they are easier for players to coordinate to. However, there are only a few known classes of such games with pure equilibria, and, to date, there has been relatively little work providing efficient and exact algorithms for computing such equilibria.

Related Work and Motivation. To this end, Monderer and Shapley [10] introduced the notion of *potential function* (and its relaxed version—(*generalised*) *ordinal potential function*) and proved that the existence of a generalised ordinal potential is equivalent to the *finite improvement property* (FIP), implying that any sequence of unilateral improving deviations terminates at a pure strategy Nash equilibrium. The authors also showed that the classes of finite—exact—potential games and congestion games coincide.

More recently, congestion games have been extended to *local-effect games* [6], *player-specific congestion games* [7], *weighted congestion games* [7] and *ID-congestion games* [9]. In these models, the player’s payoff depends not only on the number of players choosing his resources, but also on the number of players choosing the neighbor resources or on the players’ identities. Additional generalizations [11, 15, 13, 14] deal with the possibility that resources may fail to execute their assigned tasks, or with the actual order in which the tasks are executed; we will refer to this models as *congestion games with faulty or random-order services*. Finally, much of the work has been devoted to the study of the computational complexity of finding pure strategy equilibria [1, 5] and their social performance [2, 3] in the congestion games and some of extended models.

In this paper we generalise the class of congestion games still further. Specifically, we consider settings in which the payoff of a player is determined by the vector of resource congestion (note that the resources might be mutually dependent!), via any—not necessarily the sum—real-valued function. Clearly, in this very general setting a potential function and a pure strategy equilibrium are not guaranteed to exist. However, under reasonable assumptions on the structure of player strategy spaces and the payoff functions, we will prove the existence of a pure strategy equilibrium and develop a polynomial time algorithm for the computation of such an equilibrium.

Specifically, we assume that each player has a set of accessible resources, which is a subset of a given set of resources to share, and his strategy space consists of all possible subsets of his set of resources at hand. This, for instance, captures settings with typed resources, where subsets of resources of a particular type are matched to particular player tasks, or situations where some resources are not (physically) accessible by a particular player or the player’s permissions for the resource use are restricted.

In many applications of congestion games discussed in the literature, the resource utility function is *decreasing* as a function of the number of users (or, the resource cost function is increasing). This typically reflects situations where a resource is a service provider whose costs per user are increasing due to competition on internal resources, or a player’s utility from a resource decreases due to reduction in the resource efficiency or reliability caused by higher congestion. The latter also gives rise to the very real issue of redundant usage of resources, which often occurs in non-cooperative multi-agent systems, where selfish agents try to run their jobs on several resources in parallel, in an attempt to increase the probability of success or in the hope that one would be faster. In such settings it is natural to assume that the more resources a player has in use or the lower is the congestion on his utilised resources, the less marginal benefit he derives from hiring an additional resource—that is, the player’s utility is *submodular*. This is also the case in multi-task allocation settings with concave values of task portfolios and other scenarios. Finally, a player’s preference between two resources is usually determined by

the congestion these resources experience but is *independent of irrelevant alternatives*—that is, of selection and congestion levels of the other resources.

Contribution. Following the above motivation, we introduce the class of *games with congestion-averse utilities* (CAGs) where the structure of strategy spaces is the one described above, and the payoff functions are *congestion-averse*—that is, monotonically decreasing, submodular and independent of irrelevant alternatives. We observe that the presented class of games includes—but is not restricted to—the earlier discussed congestion models with player-specific payoff functions, or faulty/random-order services. Indeed, CAGs significantly generalise the aforementioned models as they, in particular:

- allow for more general—and non-identical—player payoff function structures. Informally, when utilities and costs are understood metaphorically (e.g. as a monetisation of a benefit or inconvenience) this allows us to model different degrees of motivation/impatience between players; when they are understood literally it allows us to model utility/cost differentiation on behalf of the resource provider;
- take into account the possibility that players may have unequal access to different resources. This allows us to model player-specific tasks which can only be executed by a certain collection of resources;
- allow for non-identical and mutually-dependent resources;
- can be used to model multi-task allocation and other complex scenarios.

The main results of this paper are as follows. We observe that CAGs do not, in general, admit a generalised ordinal potential function and the finite improvement property. However, we prove that every such game possesses a Nash equilibrium in pure strategies and that any strategy profile which is stable under elementary changes (adds, drops or switches of a single resource), is a Nash equilibrium—this is called *the single profitable move property* (SPMP). Moreover, we show that the family of games with SPMP coincides with the class of CAGs, and thus our result is complete. Finally, we develop a universal, polynomial time algorithm that computes a pure strategy equilibrium in any given CAG, while the methods previously developed for models with player-specific payoff functions or faulty/random-order services, which are special cases of CAGs, appear to fail in the general case. Our new technique is based on the special sequences of elementary changes which we call “drop ladders” and “swap ladders”. In particular, we show that an equilibrium can be achieved by applying $O(N^2 R^2)$ of elementary changes, where N and R stand for the number of players and resources, respectively. Most of the proofs are omitted, due to space limitations.

2 Games with congestion-averse utilities

Consider a congestion setting with a set $\mathbf{N} = \{1, \dots, N\}$ of players, where each player $i \in \mathbf{N}$ has a set \mathbf{R}_i of $R_i \in \mathbb{N}$ accessible resources, which is a subset of a finite superset $\mathbf{R} = \{r_1, \dots, r_R\}$. A player i ’s strategy is to choose a subset of resources from \mathbf{R}_i , i.e. i ’s strategy space, Σ_i , is given by a power set of \mathbf{R}_i (either including the empty set or not, depending on the nature of the particular application). We refer to a resource that a player has assigned a task to as “selected” by that player, and “unselected” otherwise.

Every N -tuple of strategies $\sigma = (\sigma_i)_{i \in \mathbf{N}}$ corresponds to an R -dimensional congestion vector $h(\sigma) = (h_r(\sigma))_{r \in \mathbf{R}}$ where $h_r(\sigma)$ is the number of players who select resource r (we drop the profile to give h_r when it's clear which profile is under consideration). Given a strategy profile $\sigma \in \Sigma$, for any player $i \in \mathbf{N}$, we define his *personalised* vector of congestion, $h^i(\sigma)$, to be a vector in \mathbb{N}^R that coincides with $h(\sigma)$ for all the resources that have been selected by i and that has zero entries for all of i 's unselected resources: that is, $h_r^i(\sigma) = h_r(\sigma)$ if $r \in \sigma_i$ and $h_r^i(\sigma) = 0$ otherwise. For a vector $h \in \mathbb{N}^R$ we define the “support” of h , $S(h) \subseteq \{1, \dots, R\}$, to be $\{j : h_{r_j} > 0\}$.

In a classic congestion game [16], the payoff function of player i is defined by $U_i(\sigma) = \sum_{r \in \sigma_i} u_r(h_r(\sigma))$, where $u_r : \{1, \dots, N\} \rightarrow \mathbb{R}$, $r \in \mathbf{R}$, is an assignment of resource utility functions; for any resource $r \in \mathbf{R}$, $u_r(k)$ denotes the utility for a player from using resource r if the total number of users of r is k .

In our, generalised, model, the utility of player i in a congestion setting is given by a function $U_i : \mathbb{N}^R \rightarrow \mathbb{R}$ that assigns a real value to a **vector** of congestion.³ To precisely define the set of utility functions we permit, it is first necessary to introduce a set of strategy modifications which we call “elementary changes”⁴. Given a profile $\sigma \in \Sigma$, we denote the elementary changes as follows:

- **Add** $A_i(r)$ —player i *adds* an unselected resource r : $\sigma'_i = \sigma_i \cup \{r\}$.
- **Drop** $D_i(r)$ —player i *drops* a selected resource r : $\sigma'_i = \sigma_i \setminus \{r\}$.
- **Switch** $S_i(r_+ \leftarrow r_-)$ —player i *switches* resources by adding resource r_+ and dropping resource r_- (note that $S_i(r_+ \leftarrow r_-) = A_i(r_+) + D_i(r_-)$)⁵

Using the above notation, we now define the “congestion-averseness” conditions on the players’ utilities. Here, a utility function is said to be congestion-averse if it (i) monotonically decreases as congestion increases, (ii) is submodular in that the “better” collection of resources a player uses—the less incentive he has to add new resources, and (iii) is independent of irrelevant alternatives (i.e., a player’s preference between two resources depends only on congestion on the resources in question). Formally,

Definition 1. A utility function $U : \mathbb{N}^R \rightarrow \mathbb{R}$ is described as **congestion-averse** if it satisfies the following three conditions:

- **Monotonicity** (M). Function U is monotonically decreasing with respect to increasing congestion: if $S(h) = S(h')$ and $\forall r, h_r \geq h'_r$, then $U(h) \leq U(h')$.
- **Submodularity** (SM). Improving a resource selection by either (i) profitable switches, (ii) extending the set of utilised resources or (iii) reducing congestion on them does not make new adds more profitable, or drops less profitable; likewise, unprofitable switches, deleting the resources or increasing the congestion does not make drops

³ Note that the player’s utility only depends on the numbers of players choosing each resource but not on their identities—that is, this setting is *anonymous* (see [4] for results on approximating equilibria in anonymous games).

⁴ In [12, 13] they are referred to as “single moves”.

⁵ Here and in what follows, “+” should be understood to mean sequential execution, read left-to-right. We also use this notation to indicate elementary changes applied to strategy profiles: e.g., $\sigma + D$ denotes a drop applied to profile σ .

more profitable, or adds less profitable. Equivalently, for any h, h' and h'' such that $|S(h)| = 1$ and $S(h) \not\subseteq S(h'), S(h'')$,

$$U(h + h') - U(h') \leq U(h + h'') - U(h''),$$

if either (i) $|S(h') \setminus S(h'')| = |S(h'') \setminus S(h')| = 1$ and $U(h') \geq U(h'')$, (ii) $S(h'') \subseteq S(h')$ and $h_j'' = h_j'$ for any $j \in S(h'')$, or (iii) $S(h') = S(h'')$ and $h' \leq h''$.

- **Independence of irrelevant alternatives** (IIA). If a player “prefers” one resource over another at their current congestion levels, then he still does so no matter what other changes are made to any other resources. Formally, if $S_i(r_+ \leftarrow r_-)$ is a profitable switch for player i given profile σ , then it is profitable for i from any other profile σ' satisfying $r_- \in \sigma'_i$, $r_+ \notin \sigma'_i$, $h_{r_-}(\sigma) = h_{r_-}(\sigma')$ and $h_{r_+}(\sigma) = h_{r_+}(\sigma')$.

A CAG is now defined as a game in the congestion domain with congestion-averse utility functions, in which a player’s utility from a combination of strategies is determined by his personalised vector of congestion. More precisely,

Definition 2. A **CAG** $\Gamma = (\mathbf{N}, \mathbf{R}, (U_i(\cdot))_{i \in \mathbf{N}})$ consists of a set \mathbf{N} of $N \in \mathbb{N}$ players, a set \mathbf{R} of $R \in \mathbb{N}$ resources, and for each player i a set of accessible resources $\mathbf{R}_i \subseteq \mathbf{R}$ and a congestion-averse utility function $U_i : \mathbb{N}^{\mathbf{R}} \rightarrow \mathbb{R}$. The strategy space for each player $i \in \mathbf{N}$ is the set of subsets of \mathbf{R}_i , and the payoff to the player from a combination of strategies σ is $u_i(\sigma) = U_i(h^i(\sigma))$, where $h^i(\sigma)$ is i ’s personalised vector of congestion as determined by σ .

As we have previously discussed, congestion-averseness is a very reasonable assumption that is natural in many applications of congestion settings. In particular, we note that the independence of irrelevant alternatives holds for classic congestion games and **all** their known up to date extensions and generalisations. Interestingly, these are the only conditions we will need on the players’ utility functions in order to guarantee a pure strategy equilibrium. Before we proceed to the proof, however, we point out some interesting subclasses of CAGs.

Interesting Subclasses. Although congestion games as a whole are unlikely to be included in the class of CAGs (as they are, in general, PLS-complete [5]), the presented model captures various scenarios that lie far beyond the borders of the classic model. Below we provide a number of examples of generalised congestion settings, previously studied in the literature, to illustrate the richness of our model. We note, however, that the class of CAGs is **not** restricted to these special cases, as we shall also learn from their properties discussed in the sequel.

Probably the first significant extension of congestion games was the work of Milchtaich on *congestion games with player-specific functions* [7], in which the utility function associated with each resource is not universal but player-specific. This generalisation was accompanied by two limiting assumptions: (i) that each player chooses only one resource, and (ii) that the utility he derives from a particular resource decreases with congestion on it. Now, one can observe that a player-specific congestion game can be easily modified to a CAG by assuming that a player can choose any (non-empty) subset of originally available single resources, and by setting the utility of a player i from a congestion vector h to be given by $U_i(h) = \sum_{r \in S(h)} u_r^i(h_r) - M(|h| - 1)$, where $u_r^i(\cdot)$ represent

resource utility functions in the original game and M is a sufficiently big number (say, $M > \sum_{r \in \mathbf{R}} \max_{i \in \mathbf{N}} u_r^i(1)$) (note that for a vector h with $S(h) = 1$ this coincides with the original resource utility function). Obviously, the sets of outcomes of these games are identical, and the congestion-averse conditions are satisfied.

Another interesting example is the family of congestion models with faulty or random-order services, that, specifically, includes *taxed congestion games with failures* [15], *congestion games with load-dependent failures* [13] and *random order congestion games*⁶ [14]. In each of these, a player has a task that can be carried out by any element of a set of independent resources. A player may decide to assign his task, simultaneously, to several resources, either for reliability reasons or hoping that his task will be completed in a short time by at least one of the resources (all possible subsets of all given resources are available to all of the players). Doing this, each player wants to maximise the probability of successful (or, quick) completion of his task and, simultaneously, to minimise his cost. It has been (naturally) assumed that the failure probabilities increase with congestion, implying the monotonicity of player utility functions, and that the marginal benefit from hiring an additional resource decreases as the player’s selection of resources improves, implying the submodularity. Finally, since resources are uncorrelated, the IIA also holds.⁷

We also note that CAGs can be used to model more complex—rather than “single job”—scenarios. Consider, e.g., a setting where a player is given different tasks, each associated with a value and workload. A subset of resources can complete a task if its productivity (which is a function of the number and congestion on the resources) meets the task’s workload. A player strives to maximise the total value of his completed tasks, and thus he is interested in executing as many (valuable) tasks as possible.

Finally, as we discussed earlier in the Introduction, CAGs possess additional features allowing for modeling player-specific tasks, non-identical and mutually dependent resources. Thus, while the above are interesting models of specific task allocation problems, the games with congestion-averse utilities provide a general framework for more realistic modelling of congestion scenarios.

3 Properties of CAGs

In this section we investigate the properties of the games with congestion-averse utilities. In particular, we observe that these games do not admit a generalised ordinal potential function and the FIP. However, as we show below, they do possess the single profitable move property. Based on this, we develop our “drop and swap ladders” technique that enables us to achieve a pure strategy Nash equilibrium in any given CAG, while the algorithms previously developed for specific subclasses fail in the general case.

3.1 The non-existence of the FIP

The finite improvement property is equivalent to the existence of a generalised ordinal potential function—a real-valued function over the set of pure strategy profiles with

⁶ These games are also referred to as *asynchronous* congestion games [12].

⁷ We omit the formal proofs and definitions for brevity of exposition.

the property that increase in the utility of a player who unilaterally shifts to another strategy implies increase in the potential function; that is, the potential increases along any improvement path.

Based on this, we construct an example showing that CAGs, in general, have no FIP. In fact, this can be concluded directly from the inclusion in the class of CAGs of congestion games with player-specific payoff functions [7], for which examples of games without FIP have been previously found. In contrast, although the models with faulty/random-order services [11, 15, 13, 14] have been shown to not admit an exact potential function, the previous work failed to prove or disprove the existence of an ordinal potential function in these games.

Note, however, that the absence of the FIP, in general, does not contradict the existence of an equilibrium in pure strategies or the convergence of particular one-sided better reply dynamics. In what follows, we consider special types of improving deviations and the corresponding properties of games with congestion-averse utilities that we will use to develop efficient procedures for constructing pure strategy equilibria.

3.2 The single profitable move property

The simplest deviations from a strategy profile in a CAG involve adds, drops or switches of single resources, referred to as *elementary changes* (see 2 for the formal definition). As we show below, any profile which is stable against elementary changes possesses no profitable deviations at all, and hence is a Nash equilibrium. The above property, referred to as the *single profitable move property* (SPMP), has been previously shown to hold for special cases of games with faulty/random-order services [15, 13, 14].⁸ Here we extend this result to the superclass of CAGs. We then complete the result by showing that the congestion-averse conditions are not just sufficient but also necessary for the existence of SPMP (see 3.4).

Theorem 1. *A strategy profile σ of a given CAG is a pure strategy Nash equilibrium if and only if it possesses no profitable elementary changes.*

This allows us to significantly reduce the size of the set of possible player deviations from a given strategy profile that we need to examine. Moreover, it is, in fact, not necessary to consider *all* adds, drops or switches, but only the *maximally* profitable ones.

Corollary 1. *A profile σ of a given CAG is a pure strategy Nash equilibrium if and only if there are no maximally profitable elementary changes available.*

We describe a strategy profile σ as *A-stable* (*D-stable*, *S-stable*) if it admits no maximally profitable adds (drops, switches); likewise for *AS-stable*, *DS-stable* and so on. Thus, another way of stating Corollary 1 is that a profile is in equilibrium if and only if it is ADS-stable.

⁸ For the class of congestion games with player-specific payoff functions [7] and, in fact, for any model with singleton strategies, the property holds trivially, as the set of all possible deviations is restricted to switches only.

3.3 Pure strategy Nash equilibrium

The SPMP has been used as a base for the proof of existence of a pure strategy Nash equilibrium in the aforementioned special cases of the games with congestion-averse utilities. However, as we shall see, the previous techniques have been heavily built on additional properties of the particular models that do not hold for more general CAGs.

Specifically, for the games with player-specific payoff functions [7] the proof is based on showing the existence of a best-reply improvement path that connects an arbitrary initial point to a Nash equilibrium. Now, any such path consists of profitable switches only (as in these games only singleton strategies are allowed), which make it impossible to use analogous dynamics in general CAGs where players have strategies of different sizes available to them, and thus a player might need to use adds or drops to modify one strategy to another, rather than just switches. In addition, an upper bound to the length of the shortest path as above provided in [7] is not polynomial. Ackermann et al [1] extended this study to deal with player-specific *matroid congestion games*, in which the strategy space of a player consists of the bases of a matroid on the set of resources, and developed a polynomial time algorithm for computing pure equilibria in such games. However, similarly to the singleton case, this dynamics involves switch-type deviations only, as each player strategies are of the same size.

In taxed congestion games with failures [15], for which a polynomial time procedure was developed, it was based on the specific property that an add applied to any DS-stable strategy profile, either preserves its DS-stability or requires only one, uniquely defined, drop for stabilisation. This property, however, does not hold even for other models with faulty/random-order services, where an add may cause a long chain of consecutive drops. Therefore, in games with load-dependent failures [13] and random order congestion games [14] the algorithms were developed to first find an initial DS-stable strategy profile whose stability is not affected by adds, termed “post-addition DS-stable profile”. In general CAGs, however, such a profile does not necessarily exist. In addition, the above algorithms only deal with identical resources and specific utility functions. This implies the need to develop a new, universal technique for computation of equilibria in games with congestion-averse utilities.

We now proceed to investigate the properties of particular sequences of elementary changes, which we call “drop and swap ladders”, when applied to partially stable strategy profiles. As we shall see, these ladders play a central role in our general method of constructing pure strategy Nash equilibria.

Drop and Swap Ladders Suppose that a strategy profile σ is AS-stable, but does have a sequence of profitable deviations, consisting of a drop followed by $m \geq 1$ switches. We define a *drop ladder* to be a sequence as above, all of whose elementary changes are maximally profitable to each of the deviators.

Definition 3. A *drop ladder* is a sequence

$$D_{i_0}(r_0) + S_{i_1}(r'_1 \leftarrow r_1) + \cdots + S_{i_m}(r'_m \leftarrow r_m),$$

consisting of a maximally profitable drop followed by a sequence of $m \geq 1$ maximally profitable switches. The **length** of the ladder is determined by the number of switches, m , and its **tail** is the last switched-out resource involved, r_m .

When applied on partially stable profiles, drop ladders have a particular structure and some interesting properties that we summarise in the following lemma.

Lemma 1. *Given a CAG, let σ be an AS-stable strategy profile that possesses a drop ladder of length m , and let σ^k denote the result of applying the drop and the first k switches to σ . Suppose further that σ^k is A-stable for $1 \leq k < m$. Then, for all $1 \leq k \leq m$, the following holds:*

- switches “chain” with one another and with the initial drop: $r'_k = r_{k-1}$;
- if there is a profitable add $A_i(r_+)$ to the profile σ^m , then $r_+ = r_m$.

Following the observation made in Lemma 1, we now define an additional class of strategy profile modifications, termed *swap ladders*, as follows.

Definition 4. *A **swap ladder** is a drop-ladder followed by a maximally profitable add at the end:*

$$D_{i_0}(r_0) + S_{i_1}(r_0 \leftarrow r_1) + \cdots + S_{i_m}(r_{m-1} \leftarrow r_m) + A_{i_{m+1}}(r_m). \quad (1)$$

*The number of switches, m , is the **length** of the ladder. The swap ladder is described as **minimal** if all intermediate strategy profiles before the last add were A-stable (i.e., if the add is performed at the first opportunity).*

By Lemma 1, a profitable add can be made only to the tail of a minimal drop ladder, and the result of the corresponding swap ladder possesses the same congestion vector as the original profile; this gives us a reason to hope that minimal swap ladders preserve AS-stability. To build the intuition for the proof, we first make a couple of observations.

Consider a swap ladder as in (1); let σ^k be the result of applying to σ the drop and the first k switches, and let σ^{m+1} be the final profile after the add. Notice that for any player i_k , $1 \leq k \leq m+1$, who performs the k 'th move after the initial drop, the congestion on his selected resources immediately before the move, i.e. at σ^{k-1} , is the same as at σ : indeed, the only resource at which congestion is any different from that at σ is r_{k-1} , which is at that point is not selected by i_k or he would be unable to switch to or add it. Likewise, after the k 'th move, at σ^{k+1} , player i_k still does not use any resource whose congestion is lower than at σ (this is since there can be only one such resource, r_k , just switched-out by i_k).

This is the key observation: that within a swap ladder, a player making an elementary change experiences the same congestion immediately before and after the move that he did before the start of the swap ladder. Swap ladders do not change congestion, so in fact this congestion is the same throughout any sequence of swap ladders. More precisely, although congestion does of course change as other players move, the congestion experienced by a given player before and after any move that this particular player makes is the same as in the initial profile, so from his decision-making point of view there is a fixed ranking on resources throughout the sequence of swap ladders. That is, for any particular congestion vector, each player has a ranking on resources determined by the utility of holding that single resource: we say that for player i , $r_1 \leq r_2$, if $U_i(\{r_1\}) \leq U_i(\{r_2\})$. The IIA property then implies that this preference is independent of what other resources the player may have (so long as congestion on r_1 and r_2 does not change): for any $x_i \in \Sigma_i$ such that $r_1, r_2 \notin x_i$, $U_i(x_i \cup \{r_1\}) \leq U_i(x_i \cup \{r_2\})$.

We are now in a position to present the key lemma which is central to our existence proof. It characterises a possible sequence of elementary changes that a given player can make in consecutive minimal swap ladders.

Lemma 2. *Consider the sequence of adds, drops and switches that a single player makes in a sequence of minimal swap ladders. Then,*

- *the resources dropped or switched-out are always the lowest-ranked among the player's selected resources right before the corresponding move,*
- *they form an increasing sequence with respect to the total rank, and*
- *once dropped or switched-out, they are not subsequently added back or switched-in.*

Lemma 2 easily supplies us with a linear bound on the number of changes that a sequence of minimal swap ladders can contain:

Corollary 2. *There can be no more than $2NR$ elementary changes in total in any sequence of minimal swap ladders.*

Proof. Consider a single player's contribution to the sequence of swap ladders. From Lemma 2, once it has been dropped or switched-out, each resource cannot subsequently be added or switched-in; each resource can therefore only be dropped or switched-out once and added or switched-in once. It follows that the total number of elementary changes for a given player is at most $2R$. The result then holds simply by multiplying by the number of players.

Finally, Lemma 2 implies the AS-stability of post-swap-ladder profiles.

Proposition 1. *If a swap ladder is applied to an AS-stable profile σ , then the resulting profile is also AS-stable.*

Based on Proposition 1 and Corollary 2 we then conclude the existence of pure strategy Nash equilibria in CAGs. The following theorem is one of our main contributions.

Theorem 2. *Every CAG possesses a pure strategy Nash equilibrium.*

3.4 Necessity of congestion-averseness conditions

As the congestion-averse conditions on utility functions have been shown to be sufficient to prove the existence of SPMP and pure strategy Nash equilibria, an interesting question that now arises is that of necessity of the above conditions. Below we show that each of the three congestion-averse conditions is necessary, in general, for the existence of SPMP.

Theorem 3. *In a CAG setting, if any one of the congestion-averse conditions on utility functions is violated then the SPMP is not guaranteed to exist.*

4 Computation of equilibria

We finally make practical use of our theoretical results. The proof of Theorem 2 suggests a constructive algorithm for finding equilibria; we can, starting from any AS-stable profile, look for maximal drop ladders, and convert them into swap ladders whenever the result is not A-stable. Obviously this process must terminate since either the total congestion strictly decreases, or we have a swap ladder, of which—courtesy of Corollary 2—there can only be a limited number consecutively. This algorithm is presented in Algorithm 1.

Algorithm 1 An algorithm for constructing an equilibrium strategy profile, which searches for drop-ladders and swap-ladders in the same loop.

```

1: INPUT: A CAG game,  $G$ 
2: OUTPUT: A pure-strategy equilibrium  $\sigma$  for  $G$ 
3:  $\sigma \leftarrow (\mathbf{R}_1, \dots, \mathbf{R}_N)$ 
4: while  $\sigma$  is not ASD-stable do
5:   choose a maximally profitable drop  $D$ 
6:    $\sigma \leftarrow \sigma + D$ 
7:   while  $\sigma$  is not AS-stable do
8:     find a maximally profitable switch  $S$ 
9:      $\sigma \leftarrow \sigma + S$ 
10:    if  $\sigma$  has a maximally profitable add  $A$  then
11:       $\sigma \leftarrow \sigma + A$ 
12:    end if
13:  end while
14: end while

```

The analysis of the worst-case asymptotic complexity of Algorithm 1 results in the following proposition.

Proposition 2. *Algorithm 1 requires $O(N^2 R^2)$ of elementary changes, and has asymptotic complexity $O(N^2 R^2 g(N, R))$, where $g(N, R)$ is the complexity of a player's utility evaluation.*

5 Conclusions

In this paper we have proved the existence of pure strategy Nash equilibria for a large class of games—Games with Congestion-Averse Utilities—loosely modelled on traditional congestion games. We have also provided an algorithm that constructs an equilibrium explicitly. This work was motivated by a desire to address a broader class of resource contention scenarios than those previously modelled, and we have indeed done so; but more remains to be done. In particular, the question of necessity of the congestion-averseness conditions for existence of pure equilibria remains open. This implies a possibility of extending our results to models without the SPMP, which is a great challenge as our current techniques build heavily on this property.

References

1. H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. In *WINE-06*, 2006.
2. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. *Lecture Notes in Computer Science*, 3884:218–229, March 2006.
3. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC-05*, pages 67–73, 2005.
4. C. Daskalakis and C. Papadimitriou. Computing equilibria in anonymous games. In *FOCS-07*, pages 83–93, 2007.
5. A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *STOC-04*, pages 604–612, 2004.
6. K. Leyton-Brown and M. Tennenholtz. Local-effect games. In *IJCAI-03*, 2003.
7. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13:111–124, 1996.
8. I. Milchtaich. Congestion models of competition. *American Naturalist*, 147(5):760–783, 1996.
9. D. Monderer. Solution-based congestion games. *Advances in Math. Economics*, 8:397–407, 2006.
10. D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
11. M. Penn, M. Polukarov, and M. Tennenholtz. Congestion games with failures. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC-05)*, pages 259–268, 2005.
12. M. Penn, M. Polukarov, and M. Tennenholtz. Asynchronous congestion games. In *AAMAS-08*, pages 1605–1608, May 2008.
13. M. Penn, M. Polukarov, and M. Tennenholtz. Congestion games with load-dependent failures: Identical resources. *Games and Economic Behavior*, to appear.
14. M. Penn, M. Polukarov, and M. Tennenholtz. Random order congestion games. *Mathematics of Operations Research*, to appear.
15. M. Penn, M. Polukarov, and M. Tennenholtz. Taxed congestion games with failures. *Annals of Mathematics and Artificial Intelligence*, to appear.
16. R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
17. T. Roughgarden and E. Tardos. How bad is selfish routing. *Journal of the ACM*, 49(2):236–259, 2002.