On the Existence of Pure Strategy Nash Equilibria in Integer–Splittable Weighted Congestion Games

Long Tran-Thanh[1], Maria Polukarov[1], Archie Chapman[2], Alex Rogers[1], and Nicholas R. Jennings[1]

[1] School of Electronics and Computer Science, University of Southampton, UK
{ltt08r,mp3,acr,nrj}@ecs.soton.ac.uk
[2] The University of Sydney Business School, Sydney, Australia
a.chapman@econ.usyd.edu.au

Abstract. We study the existence of pure strategy Nash equilibria (PSNE) in *integer–splittable weighted congestion games* (ISWCGs), where agents can strategically assign different amounts of demand to different resources, but must distribute this demand in fixed-size parts. Such scenarios arise in a wide range of application domains, including job scheduling and network routing, where agents have to allocate multiple tasks and can assign a number of tasks to a particular selected resource. Specifically, in an ISWCG, an agent has a certain total demand (aka weight) that it needs to satisfy, and can do so by requesting one or more integer units of each resource from an element of a given collection of feasible subsets.[1] Each resource is associated with a unit–cost function of its level of congestion; as such, the cost to an agent for using a particular resource is the product of the resource unit–cost and the number of units the agent requests.

While general ISWCGs do not admit PSNE (Rosenthal, 1973b), the restricted subclass of these games with linear unit–cost functions has been shown to possess a potential function (Meyers, 2006), and hence, PSNE. However, the linearity of costs may not be necessary for the existence of equilibria in pure strategies. Thus, in this paper we prove that PSNE always exist for a larger class of convex and monotonically increasing unit–costs. On the other hand, our result is accompanied by a limiting asumption on the structure of agents' strategy sets: specifically, each agent is associated with its set of accessible resources, and can distribute its demand across any subset of these resources.

Importantly, we show that neither monotonicity nor convexity on its own guarantees this result. Moreover, we give a counterexample with monotone and semi–convex cost functions, thus distinguishing ISWCGs from the class of infinitely–splittable congestion games for which the conditions of monotonicity and semi–convexity have been shown to be sufficient for PSNE existence (Rosen, 1965). Furthermore, we demonstrate that the finite improvement path property (FIP) does not hold for convex increasing ISWCGs. Thus, in contrast to the case with linear costs, a potential function argument cannot be used to prove our result. Instead, we provide a procedure that converges to an equilibrium from an arbitrary initial strategy profile, and in doing so show that ISWCGs with convex increasing unit–cost functions are weakly acyclic.

[1] Additionally, strategy sets are restricted by certain domain–specific constraints—for instance, in network routing, an agent's strategy must define a feasible flow between its given pair of source and target nodes.

G. Persiano (Ed.): SAGT 2011, LNCS 6982, pp. 236–253, 2011.

1 Introduction

The study of interaction of multiple self–interested parties ("agents") sharing commonly–available facilities ("resources") is central to computational game theory. Such settings naturally arise in a wide range of typical application domains, from traffic routing in networks (e.g. roads, air traffic or information and communication networks (Rosenthal, 1973a; Roughgarden and Tardos, 2002)), to competition in job scheduling problems (e.g. for computational services or machine scheduling (Koutsoupias and Papadimitriou, 1999)).

In many real–world scenarios in these domains, agents may find it beneficial to assign different amounts of demand to different resources, but may have restrictions on the size of units in which this demand is distributed. For example, consider a job scheduling problem, comprised of n agents and m independent machines, where each agent has several indivisable jobs to be executed. To each selected machine, an agent pays a usage cost, which is equal to the number of jobs the agent allocates to that machine multiplied by the unit–cost per job, typically depending on the total level of demand on the machine (i.e., its congestion). A similar situation arises in communication networks (e.g. the Internet), where agents send packets (or, messages) and have to decide how many packets to route on each path in the network to minimise possible delays. Additional examples for a problem of this kind may include procuring factor inputs for manufacturing processes or purchasing transport capacity for logistics networks. Importantly, in all these situations, the agents cannot split their demands in arbitrary ways, but must do so in *integer* units.

Problems of this kind are addressed in the literature as *integer–splittable weighted congestion games* (ISWCGs), where agents strategically choose from a common set of resources, and are allowed to assign multi–unit requests to each of their selected resources; however, they are constrained to make their allocations in fixed-size parts (particularly, integer units). Each resource is equipped with a "unit–cost function" that indicates the cost that each agent pays per unit of request, depending on the aggregate level of congestion on that resource (i.e., the total number of units the users contribute to the resource). Since the agents may have different congestion impacts, the cost each agent has to pay for the use of a particular resource is the product of the amount of units it requests from that resource and the corresponding unit–cost. For example, in a computational services setting, if an agent were to purchase four units of processor time from a particular service provider, it would pay the same unit–cost for all four units, with the unit–cost determined by the total demand from all agents for that resource. The overall agent's cost is given by the sum of its costs for each resource it uses. In a ISWCG, each agent has a certain integer demand (or, weight) for resource units it needs to satisfy, and its aim is to minimise the total cost of the units by distributing its weight across the available resources. Unit–cost functions are resource–specific, but are the same for all agents (i.e., resource providers cannot discriminate between users), while demands for resource units can vary across the agents. Note that the above examples are captured in the ISWCG model by identifying the set of resources with the set of machines or network links, respectively, where differences in their technical parameters and performance factors, such as efficiency, or speed, are reflected by resource–dependent costs per unit (e.g. job, or data packet). An agent's demand represents the amount of resources (job, or data traffic) each agent has, and the set of feasible assignments (task allocations, traffic routes) corresponds to the set of feasible strategy profiles.

1.1 Related Work

Much of related work deals with a traditional *congestion game* model by Rosenthal (1973a), where agents have to choose from a given finite set of resources, and where the possible choices of an agent are given by the subsets of resources that satisfy its goals. The cost of a resource is determined by the total number of its users, and an agent's overall cost is given by the sum of resource costs over the set of the agent's selected resources. In a variant setting of network congestion models, agents have to choose subsets of edges on a graph forming a path from the agent's origin to destination, in order to route their demand (i.e. flow) through the network, and the cost (e.g. latency) of each edge varies with the number of agents traversing that edge.

The important property of congestion games shown by Rosenthal (1973a) is the existence of a Nash equilibrium in pure strategies (PSNE)—a profile where each agent plays a certain (non–randomised) strategy and no one has an incentive to unilaterally change it. Such solutions are highly desirable, since, from a system–wide perspective, they imply that a system has a deterministic stable state. This is necessary in a range of control problems where randomised strategies are not appropriate (e.g. in industrial processing or transport applications). Also, unlike mixed strategy and correlated equilibria, PSNE do not rely on the assumption that agents have the capacity to accurately randomise their actions according to an equilibrium prescription.

Moreover, congestion games are also known to possess a stronger charateristic, called the "finite improvement path property" (FIP), implying that any sequence of unilateral improvement deviations (i.e., strategy changes that decrease an agent's total cost) will converge to a PSNE in finite time. This is implied by the existence of a "potential function" that decreases along any such improvement path (Monderer and Shapley, 1996). Given this, the players can use a variety of simple potential–based search processes to find a PSNE in a distributed fashion, such as fictitious play or weighted regret monitoring (Leslie and Collins, 2006; Marden *et al.*, 2007).

The traditional model has been generalised to a variety of related situations. Such generalisations, for example, include *player–specific congestion games* (Milchtaich, 1996) where an agent's payoff depends on its identity, *weighted congestion games* (Milchtaich, 1996), in which agents may have different (although fixed) congestion impacts (weights), *local–effect games* (Leyton-Brown and Tennenholtz, 2003) with an agent's cost for a particular resource being also affected by a congestion on its neighbouring resources, *congestion games with failures* (Penn *et al.*, 2009a) and *random-order congestion games* (Penn *et al.*, 2009b) modelling faulty or asynchronous resources, and *congestion–averse games* (Byde *et al.*, 2009; Voice *et al.*, 2009) where the agents' utilities are determined by general real–valued functions of congestion vectors. Note that in all these settings, agents are restricted to request only a single or a fixed number of units from each particular chosen resource; that is, in terms of the network congestion model, they have to *unsplittably* route their flow within the network.

At the other extreme, *infinitely–splittable* congestion game models assume that agents have divisible demand, which can be fractionally split acroos an *arbitrary* number of resources (paths), in any proportion (Orda *et al.*, 1993; Cominetti *et al.*, 2009). For this setting, a result from Rosen (1965) implies that PSNE are guaranteed to exist if resource cost functions are semi–convex[2] and monotone increasing. As an intermediate concept between splittable and unsplittable games, the model of *k–splittable* network congestion models was introduced by Beier *et al.* (2004) to capture scenarios where agents are restricted to split their demand across at most k different paths. However, the portion of the demand that an agent allocates to a single path can be fractional. Beier *et al.* (2004) showed that it is NP–hard to decide whether a PSNE exists within such settings. In addition, Shachnai and Tamir (2002); Krysta *et al.* (2003) obtained similar results for k–splittable congestion games in the job scheduling domain.

More relevant to our work is the paper by Meyers (2006) where the k–splittable model is modified so that agents are only allowed to allocate integer amounts of demand to each chosen resource (or, path). The authors showed that the restricted subclass of these games where the unit–cost functions are linear, possesses a potential function, and hence, the FIP holds and a PSNE is guaranteed to exist. For a general case, Rosenthal (1973b) gave an example of an asymmetric weighted network congestion game that does not have an equilibrium in pure strategies. More recently, Dunkel and Schulz (2008) strengthened this result by showing that the problem of deciding whether a weighted network congestion game with integer–splittable flows admits a PSNE is strongly NP–hard.

1.2 Our Contribution

In this paper, we extend positive results on the existence of a pure strategy equilibrium in integer–splittable congestion games to a larger class of unit–cost functions which are monotonically increasing and convex. From a practical point of view, this class is important as convex increasing costs occur in a wide range of application domains. Indeed, in many real–world systems, marginal costs typically increase as total demand increases (e.g. energy cost in smart grids or delay in multi–server systems). Furthermore, such systems are often regarded as overloaded, if the total demand exceeds a certain threshold. In this case, the users often have to pay extremely higher costs for using the resources (in smart grids, for example, each power plant has a finite production limit, and if the total demand exceeds the sum of these limits, additional expensive peaking plant must supply the excess). We note that our result is accompanied by a limiting asumption on the structure of agents' strategy sets. Specifically, we assume that each agent is associated with its set of accessible resources, wich is a part of a given superset, and can distribute its demand across any subset of these resources. For sake of brevity, in what follows we slightly abuse the notation and use the term ISWCG to define a game with such restricted strategy set structures.

The above assumption implies that negative results by Rosenthal (1973b) and Dunkel and Schulz (2008) do not apply to our setting. However, as we show, the existence of PSNE is still violated. Moreover, PSNE are not guaranteed to exist in games with

[2] A function $f(x)$ is *semi-convex* if $x \cdot f(x)$ is convex.

either non–monotone or non–convex unit–costs, implying the necessity of these conditions for PSNE existence. Interestingly, our examples show that even functions which are monotone and semi–convex result in games with no pure strategy equilibria, thus distinguishing between the classes of ISWCGs and infinitely–splittable congestion games.

Following this, our main result proves that a pure strategy equilibrium is guaranteed for ISWCGs with monotonically increasing and convex unit–costs. Importantly, as we show, PSNE exist in these games despite of the non–existence of a potential function and the FIP. Consequently, in contrast to the case with linear costs (Meyers, 2006), potential–based methods cannot be used for proving PSNE existence and finding such equilibria. Based on this, we provide a search algorithm that returns a PSNE of a given game in finite time. Finally, we note that our algorithm shows convergence from an arbitrary initial strategy profile, thus showing that convex increasing ISWCGs possess the weak acyclicity property (Monderer and Shapley, 1996).

The remainder of the paper unfolds as follows. First, in Section 2 we formally define the model for ISWCGs. Then, in Section 3 we show that no guarantees on PSNE existence can be made if the unit–cost functions are not convex or monotone increasing. Following this, in Section 4 we study the case of ISWCGs with convex increasing costs. We show that these games do not generally possess a potential function by giving an example of an improvement cycle. Nonetheless, we prove that they are guaranteed to possess PSNE if the cost function is convex and monotone increasing, and devise an algorithm for computing them. Due to space limitations, some of the proofs are ommited from this version of the paper.

2 The Model

Consider a congestion domain with a set $N = \{1, \ldots, n\}$ of agents, where each agent $i \in N$ has a set R_i of $m_i \in \mathbb{N}$ accessible resources, which is a subset of a finite superset $R = \{r_1, \ldots, r_m\}$. An agent i needs to execute $X^i \in \mathbb{N}$ task units, and can distribute this demand (or, *weight*) arbitrarily among its resources. Note that each agent can use more than one *integer* unit from a single selected resource. An agent i's (pure) strategy is given by $x^i = \left(x_r^i\right)_{r \in R}$, where $x_r^i \in \mathbb{N}$ is the number of units that agent i demands from resource $r \in R$, such that $x_r^i = 0$ for all $r \notin R_i$ and

$$\sum_{r \in R} x_r^i = \sum_{r \in R_i} x_r^i = X^i \tag{1}$$

Every combination of strategies (a *strategy profile*) $x = \left(x^i\right)_{i \in N}$ corresponds to a *congestion vector* $h(x) = \left(h_r(x)\right)_{r \in R}$, where

$$h_r\left(x\right) = \sum_{i \in N} x_r^i \tag{2}$$

indicates the *congestion*—the total number of assigned tasks (or, demanded units) on resource $r \in R$ in profile x.

From the perspective of agent i, a strategy profile x can be viewed as $\left(x^i, x^{-i}\right)$, where x^{-i} stands for the joint strategy of other agents. Similarly, for $r \in R$ we denote by

$$h_r^{-i}(x) = \sum_{j \neq i} x_r^j = h_r(x) - x_r^i \tag{3}$$

the congestion on resource r incurred by the collective demand of the agents, excluding agent i.

Each resource $r \in R$ is associated with a *unit–cost* (or simply, a *cost*) function $c_r : \mathbb{N} \to \mathbb{R}$ defining the cost for a unit of demand on resource r as a function of the total congestion on the resource. For simplicity, it is convenient to assume that cost functions are non–negative, although our results do not rely on this assumption.

Given this, the *payoff function* of an agent is defined as follows. The overall cost agent i has to pay in a strategy profile x is

$$C^i(x) = \sum_{r \in R} x_r^i c_r(h_r(x)) \tag{4}$$

Furthermore, the total cost of the system is given by

$$C(x) = \sum_{i \in N} C^i(x) = \sum_{r \in R} h_r(x)\, c_r(h_r(x)) \tag{5}$$

Definition 1. *An* integer–splittable weighted congestion game (ISWCG) $\Gamma = \left(N, R, \left(X^i\right)_{i \in N}, \left(c_r(\cdot)\right)_{r \in R}\right)$ *consists of a set N of $n \in \mathbb{N}$ agents, a set R of $m \in \mathbb{N}$ resources, a unit–cost function c_r for each resource, and for each agent i a set of accessible resources $R_i \subseteq R$ and a total demand (aka weight) X^i. The strategy set for each agent $i \in N$ is the set of m-dimensional vectors $\left\{\left(x_r^i\right)_{r \in R} \in \mathbb{N}^m\right\}$, such that $\sum_{r \in R} x_r^i = X^i$, $x_r^i = 0\ \forall r \notin R_i$, and the cost to the agent for a combination of strategies x is $C^i(x) = \sum_{r \in R} x_r^i c_r(h_r(x))$, where $h_r(x)$ is the vector of congestion as determined by x.*

3 Non-existence of PSNE

In this section, we show that general ISWCGs do not necessarily admit pure strategy Nash equilibria (PSNE). We provide two examples, based on which we reason about conditions that would guarantee PSNE existence.

Example 1. Consider a two–player ISWCG with demands $X_1 = 2$ and $X_2 = 1$, and two resources with the following unit–cost functions:

$$c_{r_1}(1) = 12,\ c_{r_1}(2) = 5,\ c_{r_1}(3) = 7$$
$$c_{r_2}(1) = 10,\ c_{r_2}(2) = 6,\ c_{r_2}(3) = 10$$

The payoff matrix of the game is presented in Table 1. One can easily verify that there is no PSNE in this game.

Table 1. No PSNE in ISWCGs with non-monotone unit–costs

	$(0,2)$	$(1,1)$	$(2,0)$
$(0,1)$	10, 20	6, 18	10, 10
$(1,0)$	12, 12	5, 15	7, 14

Note that the cost functions in Example 1 are not monotone, but convex. That is, the convexity condition on its own is not sufficient for the existence of a pure strategy equilibrium. The next example demonstrates that neither is monotonicity sufficient.

Example 2. Consider a two–player ISWCG with demands $X_1 = 3$ and $X_2 = 1$, and two identical resources with a unit–cost function $c_{r_1}(\cdot) = c_{r_2}(\cdot) = c_r(\cdot)$ given by:

$$c_r\,(1) = 3,\; c_r\,(2) = 8,\; c_r\,(3) = 10\, c_r\,(4) = 12$$

The payoff matrix of the game is presented in Table 2. Inspection shows that there is no PSNE in this game.

Table 2. No PSNE in ISWCGs with non-convex cost functions

	(0,3)	(1,2)	(2,1)	(3,0)
(1,0)	3, 30	8, 24	10, 23	12, 36
(0,1)	12, 36	10, 23	8, 24	3, 30

As mentioned above, Example 1 is convex, while Example 2 is monotone–increasing, implying that if either property of the cost functions is violated, a PSNE is not guaranteed. Furthermore, the cost function $c_r(x)$ in Example 2 is *semi–convex* (i.e., $x \cdot c_r(x)$ is convex). It implies that the conditions of monotonicity and semi-convexity, which have been shown to be sufficient for PSNE existence in infinitely–splittable congestion games, do not apply to the integer–splittable case! Based on this, in the following section we prove that a pure strategy equilibrium always exists in the ISWCGs whose resource unit–cost functions are monotone–increasing and convex.

4 Convex Increasing ISWCGs

In this section, we investigate the subclass of ISWCGs with convex and monotonically increasing cost functions (henceforth, *convex increasing* ISWCGs). Our main result proves that pure strategy Nash equilibria always exist in such games. Importantly, as we show in 4.1, an arbitrary sequence of myopic improving deviations may cycle even in this case; hence, the FIP property does not hold and a potential function argument is not applicable. Against this background, in 4.2 we propose a special dynamic procedure, that reaches an equilibrium from any starting point. This shows that convex increasing integer–splittable congestion games possess the weak–acyclicity property and implies an algorithm for finding PSNE in these games.

4.1 Violating the Finite Improvement Property

Given a pure strategy profile of a game, consider an arbitrary sequence of unilateral moves, where at each step a deviating agent improves its payoff with respect to the current one it gets from the game. If such a sequence of myopic improvement steps terminates, the resulting strategy profile is a Nash equilibrium. Now, if *every* such path leads to a PSNE, it is said that the game has the *finite improvement path property*

(FIP). Importantly, the FIP is equivalent to the existence of a *generalised ordinal potential* (Monderer and Shapley, 1996)—a real-valued function over the set of pure strategy profiles that strictly decreases along any improvement path. Thereby, if the FIP holds for a particular game, then the agents only need to search for a local minimum point of the potential, in order to find a pure strategy equilibrium. It is known that Rosenthal's congestion games always possess a potential function and the FIP and, in fact, are a central class of games with this property (see Monderer and Shapley (1996) for a detailed discussion).

Below, we demonstrate that convex increasing ISWCGs do not fall within the framework of congestion games, as these games generally violate the FIP property. Specifically, we provide an example of the convex increasing ISWCG that contains an improvement cycle, as follows.

Example 3. Consider a convex increasing ISWCG game with 2 agents $N = \{1, 2\}$ and 5 resources $R = \{r_1, r_2, \ldots, r_5\}$, where both agents have access to all of the resources. Agent 1 requires 14 units of resources, and agent 2's demand is 36. The unit–cost functions have the following particular values:

$$c_{r_1}(1) = 39$$
$$c_{r_2}(1) = 350$$
$$c_{r_3}(35) = 5,\ c_{r_3}(36) = 8,\ c_{r_3}(37) = 21$$
$$c_{r_4}(1) = 150$$
$$c_{r_5}(13) = 16,\ c_{r_5}(14) = 22,\ c_{r_5}(15) = 52$$

Consider profile $x = (x^1, x^2)$, where $x^1 = (1, 0, 10, 0, 3)$ and $x^2 = (0, 1, 25, 0, 10)$, with a corresponding congestion vector $h(x) = (1, 1, 35, 0, 13)$. Accordingly, the vector of unit–cost values as determined by x is $(39, 350, 5, 0, 16)$, and the agents' overall costs are $C^1(x) = 1{\cdot}39{+}10{\cdot}5{+}3{\cdot}16 = 137$ and $C^2(x) = 1{\cdot}350{+}25{\cdot}5{+}10{\cdot}16 = 635$. We construct an improvement cycle that starts at x and consists of simple improvement steps at which an agent moves a *single* task unit from one resource to another. First, agent 1 moves 1 unit from r_1 to r_3. The resulting cost to agent 1 is then given by $11 \cdot 8 + 3 \cdot 16 = 136$, which is less by 1 than what the agent paid before. Following this, agent 2 moves a unit from r_2 to r_3 and gets $26 \cdot 21 + 10 \cdot 16 = 706$, thus reducing the cost of $1 \cdot 350 + 25 \cdot 8 + 10 \cdot 16 = 710$ it paid after the first improvement step by agent 1. The whole sequence of moves and the corresponding cost reductions to deviating

Table 3. Improvement cycle in ISWCGs with convex increasing unit–cost functions

Step	**Deviator**	**Move**	**Improvement**
1	Agent 1	1 unit $r_1 \to r_3$	137 - 136 = 1
2	Agent 2	1 unit $r_2 \to r_3$	710 - 706 = 4
3	Agent 1	1 unit $r_3 \to r_4$	279 - 278 = 1
4	Agent 2	1 unit $r_3 \to r_5$	368 - 367 = 1
5	Agent 1	1 unit $r_4 \to r_5$	266 - 258 = 8
6	Agent 2	1 unit $r_5 \to r_2$	697 - 695 = 2
7	Agent 1	1 unit $r_5 \to r_1$	138 - 137 = 1

agents is listed in Table 3. Note that after 7th step the system turns back to the initial strategy profile, and so the improvement path cycles.

However, the non-existence of the FIP and a potential function in a class of games does not generally contradict the existence of an equilibrium in pure strategies. Thus, in the following section, we prove that convex increasing integer–splittable congestion games do always possess such an equilibrium, despite of the non–existence of the FIP. Our proof is constructive and yields a natural procedure that achieves an equilibrium point in a finite number of steps. Importantly, the convergence is guaranteed, regardless of the initial strategy profile, and so convex increasing congestion games with multi–unit resource demands are weakly–acyclic.

4.2 Nash Equilibria

We start with the following Lemma 1, introducing a useful property of convex increasing functions that we will employ in proving results within this section.

Lemma 1. *Let* $c : \mathbb{N} \to \mathbb{R}$ *be a convex and monotonically increasing function. Then, for any* $0 \leq x \leq y$ integer *and* $h \geq 0$*, the following holds:*

- $yc(h+y) - xc(h+x) \geq (y-x)\left[(x+1)c(h+x+1) - xc(h+x)\right]$
- $yc(h+y) - xc(h+x) \leq (y-x)\left[yc(h+y) - (y-1)c(h+y-1)\right]$

Moreover, the inequalities are strict if $y > x+1$.

We now turn to prove our main result. In doing so, we first provide a useful characterisation of best response strategies in ISWCGs with convex increasing costs (Theorem 1). We then use this characterisation to prove PSNE existence (Theorem 2) and define a special type of improvement dynamics (Algorithm 1) that converges to a Nash equilibrium from an arbitrary starting point (Theorem 3).

Distances between Strategies

Definition 2. *The* modified Hamming distance *between agent* i*'s strategies* $x^i = \left(x^i_r\right)_{r\in R}$ *and* $y^i = \left(y^i_r\right)_{r\in R}$ *is defined as*

$$H\left(x^i, y^i\right) = \sum_{r\in R}\left|x^i_r - y^i_r\right| \tag{6}$$

Now, since equation (1) must hold for any strategy of agent i, from Definition 2 we easily derive the following lemma.

Lemma 2. *In an integer–splittable congestion game, if* $x^i \neq y^i$ *are different strategies of agent* i*, then* $H\left(x^i, y^i\right) \geq 2$.

Based on this lemma, if the modified Hamming distance between two strategies x^i and y^i is exactly 2, we will refer to them as *neighbours*. The next lemma then states that an improving deviation from a particular strategy (if one exists) can always be found among its neighbours.

Lemma 3. *Let* $x = \left(x^i, x^{-i}\right)$ *be a strategy profile of a given ISWCG with convex increasing costs. If* x^i *is not agent* i*'s best response against* x^{-i}*, then there exists a strategy* y^i*, such that* $H\left(x^i, y^i\right) = 2$ *and* $C^i\left(y^i, x^{-i}\right) < C^i(x)$.

Single Unit Moves. Given this, it will be useful to identify best improving deviations within the set of neighbouring strategies.

Definition 3. *Let $D^i_{\max}(x)$ denote the value of maximal improvement that agent i can achieve by deviating to a neighbouring strategy from profile x. That is,*

$$D^i_{\max}(x) = \max_{y^i:\, H\left(x^i, y^i\right)=2} \left\{ C(x) - C\left(y^i, x^{-i}\right) \right\} \tag{7}$$

Obviously, if x is a Nash equilibrium profile, then for any $i \in N$ we have $D^i_{\max}(x) \leq 0$. Otherwise, if for some agent i its strategy x^i is not a best response against x^{-i}, then by Lemma 3, there exists a strategy y^i for agent i such that $H\left(x^i, y^i\right) = 2$ and $D^i_{\max}(x) \geq U(x) - U\left(y^i, x^{-i}\right) > 0$. This implies the following theorem.

Theorem 1. *Given a convex increasing ISWCG, a strategy x^i is a best response to agent $i \in N$ against s^{-i}* if and only if $D^i_{\max}(x) \leq 0$.

Thereby, a strategy profile x is a PSNE if and only if the condition in Theorem 1 holds for each agent $i \in N$. We seek such a profile by constructing an improvement path, where at each step an agent deviates to a best neighboring strategy. Let us now characterise these improving moves.

From Lemma 2, it is easy to see that x^i and y^i are neighboring strategies of agent $i \in N$ if and only if there are $p, q \in R_i$ such that $y^i_p = x^i_p - 1$ and $y^i_q = x^i_q + 1$. That is, agent i deviates from x^i to y^i by moving exactly one task unit from resource p to resource q. Hereafter, we refer to such deviations as *single unit moves*.

Let $D^i_{p\to q}(x)$ denote agent i's value of improvement by taking a single unit move $p \to q$ from profile x. That is,

$$D^i_{p\to q}(x) = C^i(x) - C^i\left(y^i, x^{-i}\right) \tag{8}$$

where y^i is such that $y^i_p = x^i_p - 1$, $y^i_q = x^i_q + 1$ and $y^i_r = x^i_r$ for all $r \in R \setminus \{p, q\}$. Given this, we can rewrite $D^i_{\max}(x)$ as follows:

$$D^i_{\max}(x) = \max_{p \neq q \in R_i} D^i_{p\to q}(x) \tag{9}$$

Now, let us calculate

$$\begin{aligned} D^i_{p\to q}(x) = &\left[x^i_p c_p(h_p(x)) - \left(x^i_p - 1\right) c_p(h_p(x) - 1)\right] \\ &+ \left[x^i_q c_q(h_q(x)) - \left(x^i_q + 1\right) c_q(h_q(x) + 1)\right] \end{aligned} \tag{10}$$

and consider

$$p^{i*} \in \arg \max_{r \in R_i:\, x^i_r > 0} \{x^i_r c_r(h_r(x)) - \left(x^i_r - 1\right) c_j(h_r(x) - 1)\} \tag{11}$$

That is, resource p^{i*} guarantees to agent i a maximal cost reduction if it removes one unit of demand from that resource. Similarly, resource

$$q^{i*} \in \arg \min_{r \in R_i} \{\left(x^i_r + 1\right) c_j(h_r(x) + 1) - x^i_r c_j(h_r(x))\} \tag{12}$$

guarantees a minimal increase in cost when i adds one unit of demand to q^{i*}.

Obviously, for any pair of resources p and q with $x_p^i > 0$ we have that $D_{p \to q}^i(x) \leq D_{p^{i*} \to q^{i*}}^i(x)$. That is, if $p^{i*} \neq q^{i*}$ then $D_{\max}^i(x) = D_{p^{i*} \to q^{i*}}^i(x)$, and if $D_{\max}^i(x) > 0$ then $p^{i*} \to q^{i*}$ is a *best single unit move* to agent i from x. Otherwise, if $p^{i*} = q^{i*}$, then the following lemma implies that x^i is a best response strategy to agent i.

Lemma 4. *Given a convex increasing ISWCG and a strategy profile x, if for agent $i \in N$ there exist p^{i*} and q^{i*} (as defined in equations (11) and (12), respectively) such that $p^{i*} = q^{i*}$, then $D_{\max}^i(x) \leq 0$.*

Best Response Dynamics. Let x be an arbitrary strategy profile of a given ISWCG with convex increasing costs. As we concluded before from Theorem 1, if $D_{\max}^i(x) \leq 0$ holds for every agent $i \in N$ then x is a Nash equilibrium strategy profile. So assume otherwise, and let i be an agent with $D_{\max}^i(x) > 0$. By Lemma 4, we have that $p^{i*} \neq q^{i*}$, and let $B^i(x)$ denote the number of best single unit moves of i from x. We prove the following.

Theorem 2. *Given an ISWCG with convex increasing costs, let x be a strategy profile which is not in equilibrium. Then, there exists a profile y, such that for each agent $i \in N$, one of the following three conditions is satisfied:*

1. $D_{\max}^i(x) > D_{\max}^i(y)$
2. $D_{\max}^i(x) = D_{\max}^i(y)$ *and* $B^i(x) > B^i(y)$
3. $D_{\max}^i(x) = D_{\max}^i(y)$ *and* $B^i(x) = B^i(y)$

Moreover, for at least one agent either 1. or 2. holds.

Corollary 1. *Given an ISWCG with convex increasing costs and a strategy profile x, let*

$$P(x) = L \cdot \sum_{i \in N} D_{\max}^i(x) + \sum_{i \in N} B^i(x) \qquad (13)$$

where L is a large number satisfying $L \geq \frac{nm(m-1)}{\min_{p,q,k,l} |c_p(k) - c_q(l)|}$. Then, if x is not a Nash equilibrium, then there exists a profile y, such that $P(x) > P(y)$.

Note that function $P(\cdot)$ in (13) does not decrease along *any* improvement path, and so the FIP does not follow. Nonetheless, Theorem 2 and Corollary 1 imply the existence of pure strategy Nash equilibria in convex increasing integer–splittable congestion games. To prove Theorem 2 we need the following auxiliary lemma.

Lemma 5. *Given a convex increasing ISWCG, assume there is a sequence $(x_1, x_2, \ldots x_T)$ of strategy profiles such that:*

- *x_1 is not a pure strategy Nash equilibrium, and x_2 is obtained from x_1 by a best single unit move of some agent i with $D_{\max}^i(x_1) > 0$*
- $\forall 1 < t < T, \exists r_t^+, r_t^- \in R$, *such that* $h_{r_t^+}(x_t) = h_{r_t^+}(x_1) + 1$, $h_{r_t^-}(x_t) = h_{r_t^-}(x_1) - 1$. *Furthermore,* $\forall r \in R \setminus \{r_t^+, r_t^-\}$ *we have* $h_r(x_t) = h_r(x_1)$

– *$\forall 1 < t < T$, $\exists j_t \in N$ with $D^{j_t}_{\max}(x_t) > D^{j_t}_{\max}(x_1)$ or $D^{j_t}_{\max}(x_t) = D^{j_t}_{\max}(x_1) \wedge B^{j_t}(x_t) > B^{j_t}(x_1)$, and $\exists r \in R_{j_t}$, such that either $D^{j_t}_{\max}(x_t) = D^{j}_{r_t^+ \to r}(x_t)$ or $D^{j_t}_{\max}(x_t) = D^{j_t}_{r \to r_t^-}(x_t)$. Furthermore, $x_{t+1}$is obtained from x_t by the corresponding best single unit move by agent j, that removes a unit from r_t^+ (or adds one to r_t^-). Moreover, if $r = r_t^+$ or $r = r_t^-$ (i.e., $D^{j}_{\max}(x_t) = D^{j}_{r_t^+ \to r_t^-}(x_t)$), then $t + 1 = T$.*

Then, for all $1 < t < T$ we have $D^{j_t}_{\max}(x_{t+1}) < D^{j_t}_{\max}(x_1)$ or $D^{j_t}_{\max}(x_{t+1}) = D^{j_t}_{\max}(x_1) \wedge B^{j_t}(x_{t+1}) < B^{j_t}(x_1)$.

That is, at each step t in the sequence, we have an agent j_t, whose current maximal improvement is higher than the value it had in the initial strategy profile x_1 (or the number of best single unit moves available to i at step t is greater than that it had at the first step). Furthermore, the congestion levels in x_t differ from congestion levels in the initial profile x_1 for only two resources r_t^+ and r_t^-, plus/minus one unit each. A best move of agent j_t is to either move a unit from r_t^+ to some resource r (and so $r = r_{t+1}^+$, unless $r = r_t^-$), or to take one from some r and add to r_t^- (in which case, $r = r_{t+1}^-$, unless $r = r_t^+$). This best move by j_t then results in the subsequent strategy profile x_{t+1}, and if $r = r_t^-$ or $r = r_t^+$ (i.e., agent j_t's best move is from r_t^+ to r_t^-), then this is the last move in the sequence. Now, if such a sequence exists in a given game, then at each iteration, the value of maximal improvement for the corresponding deviator (or the number of its available best single unit moves) decreases comparing to what it had in the initial point of the sequence x_1.

Proof (of Theorem 2). We construct a finite sequence of best single unit moves that results in a strategy profile y for which the theorem holds. In particular, we first prove that during the sequence, if we reach a certain congestion level profile twice, then we can leave out the in between steps. Using this result, we then show that we cannot infinitely continue the sequence without reaching a strategy profile for which the theorem holds.

In doing so, we define a particular order of moves, as follows. Let $\{x_1, x_2, \ldots\}$ denote the sequence of strategy profiles resulted from a sequence of best single unit moves $x_t \to x_{t+1}$, $t \geq 1$, as defined in Lemma 5. We refer to the moves $r_t^+ \to r$ and $r_t^- \to r$ as *forward* and *backward* moves, respectively (moves $r_t^+ \to r_t^-$ can be both, but we will make it clear in the context). Note that by Lemma 5, if at step t some agent i violates the conditions of the theorem, then there is always either a forward or a backward move that it can apply. Given this, we start the sequence with a series of forward moves, and when no such move is available, we switch to backward moves if any exist. We prove that this construction leads to a desired strategy profile in any case. The steps involved within the proof are described below.

Step 1: By definition, we move from x_1 to x_2 with some agent i who applies its best single unit move. From x_2, we only allow agents to make forward moves (if exist); that is, for now, backward moves are out of consideration. Let $\{r_{f_1}, r_{f_2}, r_{f_3}, \ldots\}$ denote the sequence of such forward moves, where $r_{f_t} \to r_{f_{t+1}}$ denotes a forward move from resource r_{f_t} to $r_{f_{t+1}}$ at step t. For the sake of simplicity, we assume that the first move, that deviates x_1 to x_2, is also a forward move (i.e., that move is $r_{f_1} \to r_{f_2}$, and we start the sequence from the initial strategy profile).

Now, consider the case where $\exists u, v, u < v$, such that $r_{f_u} = r_{f_v}$, and none of them is equal to r_{f_1}; that is, the sequence of forward moves creates a loop by turning back to a previously visited resource, which is not the first resource. We show that if the sequence is $\{r_{f_1}, r_{f_2}, \ldots, r_{f_u}, \ldots, r_{f_v}, r_{f_{v+1}}, \ldots\}$, then if agent i is the one who makes the move $r_{f_v} \to r_{f_{v+1}}$, then it can make the move $r_{f_u} \to r_{f_{v+1}}$ as well; and thus, the sequence $\{r_{f_1}, r_{f_2}, \ldots, r_{f_u}, r_{f_{v+1}}, \ldots\}$ is also a feasible sequence of forward moves. That is, we can leave out the loop $\{r_{f_{u+1}}, \ldots, r_{f_v}\}$, without violating the conditions of Lemma 5.

Let x_{f_u} and x_{f_v} denote the profiles that result from subsequences $\{r_{f_1}, r_{f_2}, \ldots, r_{f_u}\}$ and $\{r_{f_1}, r_{f_2}, \ldots, r_{f_u}, \ldots, r_{f_v}\}$, respectively. Now, suppose that agent i makes the move $r_{f_v} \to r_{f_{v+1}}$ from x_{f_v}. We show that this move is available to i at x_{f_u} as well. Indeed, consider the congestion level and the demand of agent i on resource r_{f_u} in x_{f_u} and x_{f_v}. Since all the moves are forward moves, it is easy to see that the congestion level in both profiles is given by $h_{r_{f_u}}(x_1) + 1$ (this is since $r_{f_u} = r^+_{f_u} = r^+_{f_v}$). Furthermore, it can be shown that the demand of agent i on r_{f_u} in x_{f_v} is at most as high as it was in x_{f_u}. One exceptional case is when agent i is the one who makes the move $r_{f_{v-1}} \to r_{f_v}$, and thus its demand on r_{f_u} may be increased by one. However, Lemma 5 implies that whenever an agent makes a best unit move (either forward or backward), at next step of the sequence it satisfies the conditions of the theorem. Hence, it cannot be the one who makes the subsequent move. Given this, if agent i is the one who makes the move $r_{f_{v-1}} \to r_{f_v}$, then it cannot make the move $r_{f_v} \to r_{f_{v+1}}$, which is a contradiction. This implies that the demand of agent i on r_{f_u} cannot be greater in x_{f_v} than that it has in x_{f_u}.

Now, if the demand of agent i on resource $r_{f_{v+1}}$ in x_{f_v} is smaller than its demand on the same resource in x_{f_u}, then we show that this results in a contradiction. We prove this by indirection; that is, suppose that it is true. This implies that there exists $u < z < v$, such that $r_{f_z} = r_{f_{v+1}}$, and agent i moved a unit from r_{f_z} to $r_{f_{z+1}}$ within the sequence. Furthermore, the demand of agent i on $r_{f_z} = r_{f_{v+1}}$ after the move is decreased by 1, compared to its demand in x_{f_u}. That is, since the demand of agent i on $r_{f_{v+1}}$ in x_{f_v} is smaller than in x_{f_u}, agent i must move some units from that resource in between. Thus, we focus on the first move among these, which decreases agent i's demand by 1. Note that, by definition of the sequence, $r_{f_{v+1}} \to r_{f_{z+1}}$ is a best single unit move of agent i in x_{f_z}. Now, let a denote the amount of agent i's decreased cost by removing one unit from $r_{f_{v+1}}$ in x_{f_z}, and b denote the agent's increased cost by adding one unit to $r_{f_{z+1}}$, also in x_{f_z}. Thus, the improvement that agent i gets by making $r_{f_{v+1}} \to r_{f_{z+1}}$ is $a - b > 0$. Similarly, let c denote the amount of agent i's decreased cost by removing one unit from r_{f_v} in x_{f_v}, and d denote the agent's increased cost by adding one unit to $r_{f_{v+1}}$ (i.e. r_{f_z}), also in x_{f_v}. Since $r_{f_v} \to r_{f_{v+1}}$ is also a best single move, $c - d > 0$. It is easy to see that both the congestion level and the demand of agent i on $r_{f_{v+1}}$ remain the same after the move $r_{f_v} \to r_{f_{v+1}}$, and before the move $r_{f_{v+1}} \to r_{f_{z+1}}$. Thus, we have $d = a$; and thus, $c - b > a - b > 0$. This implies that in x_{f_z}, the best single move is not moving from $r_{f_{v+1}}$ to $r_{f_{z+1}}$, but from r_{f_v} (since both the congestion level and agent i's demand on r_{f_u} is not modified between x_{f_u} and x_{f_v}; that is, it stays unchanged within the loop). This, however, is a contradiction, since $r_{f_{v+1}} \to r_{f_{z+1}}$ is supposed to be the best single move in x_{f_z}.

Given this, the demand of agent i on resource $r_{f_{v+1}}$ in x_{f_v} is at least as its demand on the same resource in x_{f_u}. In this case, $r_{f_u} \to r_{f_{v+1}}$ is feasible for agent i from x_{f_u} as well. Indeed, since $r_{f_u} \to r_{f_{v+1}}$ is feasible for agent i in x_{f_v}, such that the demand of agent i on r_{f_u} in x_{f_v} is not higher than in x_{f_u}, and its demand on $r_{f_{v+1}}$ in x_{f_v} is not smaller than in x_{f_u}. That is, by choosing $r_{f_w} = r_{f_{v+1}}$, we get that $\{r_{f_1}, r_{f_2}, \ldots, r_{f_u}, r_{f_w}, \ldots\}$ is also a feasible sequence, without the $\{r_{f_{u+1}}, \ldots, r_{f_v}\}$ loop.

Thus, in summary, we can say that if there's a loop within the sequence, that does not return to r_{f_1}, then we can leave that loop out of the sequence.

Step 2: Now, we will show that if the sequence does not return to r_{f_1}, then it has to be finite. We prove this by contradiction as follows: Suppose that the sequence is infinite and never returns to r_{f_1}. Given this, there is an infinite subsequence of moves $r_{f_{u(t)}} \to r_{f_{u(t)+1}}$ applied by a particular agent i, such that $r_{f_{u(1)}} = r_{f_{u(2)}} = \ldots$ and $r_{f_{u(1)+1}} = r_{f_{u(2)+1}} = \ldots$. That is, agent i makes the same move $r_{f_{u(t)}} \to r_{f_{u(t)+1}}$ infinitely many times within the sequence. Furthermore, the demand of i on resources $r_{f_{u(t)}}$ and $r_{f_{u(t)+1}}$ are the same for every t. That is, if agent i's demands on $r_{f_{u(1)}}$ and $r_{f_{u(1)+1}}$ are a and b, respectively, then they are a and b for any t.

Now, consider the move $r_{f_{u(1)}} \to r_{f_{u(1)+1}}$ of agent i. After this move, agent i's demand on $r_{f_{u(1)}}$ and $r_{f_{u(1)+1}}$ becomes $a-1$ and $b+1$, respectively. However we know that when agent i makes the move $r_{f_{u(2)}} \to r_{f_{u(2)+1}}$, these values return to a and b again. That is, before applying $r_{f_{u(2)}} \to r_{f_{u(2)+1}}$, agent i had to make a move $r_{f_v} \to r_{f_{v+1}}$, where $r_{f_{v+1}} = r_{f_{u(1)}} = r_{f_{u(2)}}$, to increase its demand on $r_{f_{u(2)}}$ back to a. Now, note that $u(1) < v < u(2) - 1$. This implies that the subsequence $\{r_{f_{v+2}}, \ldots, r_{f_{u(2)}}\}$ forms a loop, and thus, according to the claim we stated in Step 1, we can leave this loop out from the sequence. That is, the moves $r_{f_v} \to r_{f_{v+1}}$ and $r_{f_{u(2)}} \to r_{f_{u(2)+1}}$ become subsequent moves within the sequence. However, as Lemma 5 implies, none of the agents can subsequently make more than one move within the sequence, and thus, this situation is not possible. This contradicts the initial assumption, and hence, sequence $\{r_{f_1}, r_{f_2}, \ldots\}$ either returns to r_{f_1}, or it is finite.

Step 3: Based on the results described in Step 2, if $\{r_{f_1}, r_{f_2}, \ldots\}$ (i.e. the sequence of forward moves) is not finite, then it has to return to r_{f_1}. That is, $\exists v$ such that in $\{r_{f_1}, r_{f_2}, \ldots, r_{f_v}\}$, $r_{f_v} = r_{f_1}$. If there is an inner loop within this sequence, then we can remove that loop (as proved in Step 1). Thus, we can assume that the sequence does not contain any inner loops (note that the sequence itself is also a loop). Let x_{f_v} denote the resulting strategy profile by making this sequence of forward moves. We show below that x_{f_v} satisfies the conditions of the theorem; that is, it is the strategy profile we are looking for.

Note that by returning to r_{f_1}, the congestion level on all the resources in x_{f_v} is the same as it is in x_1. Since the sequence does not contain any inner loops, it is easy to see that for any agent i, there is a set of disjoint pairs of resources $r_{f_{u(k)}}, r_{f_{u(k)+1}}$ such that agent i makes the move $r_{f_{u(k)}} \to r_{f_{u(k)+1}}$ within the sequence. This indicates that in x_{f_v}, agent i's demand on $r_{f_{u(k)}}$ is decreased by 1, compared to that it has on that resource in x_1 (since agent i removes one unit from that resource). On the other hand, agent i's demand on $r_{f_{u(k)+1}}$ is increased by 1, compared to that it has on that resource in x_1.

In order to prove the claim above, we show that the value of a best unit move of agent i in x_{f_v} is decreased, compared to that it has in x_1 (or the number of such moves is decreased). Since the congestion level is the same on all the resources in the two strategy profiles, we just need to consider the cases where agent i makes a move from $r_{f_{u(k)+1}}$ (where the demand is increased) to $r_{f_{u(l)}}$ (where the demand is decreased) for a particular pair of k, l.

If $k = l$, then $r_{f_{u(k)}} \rightarrow r_{f_{u(k)+1}}$ is a forward move of agent i. Let $x_{f_{u(k)}}$ and $x_{f_{u(k)+1}}$ denote the strategy profiles before and after the move. If $r_{f_{u(k)}} \rightarrow r_{f_{u(k)+1}}$ is not the first move in the sequence, then the congestion levels on resources $r_{f_{u(k)}}$ and $r_{f_{u(k)+1}}$ in $x_{f_{u(k)}}$ and $x_{f_{u(k)+1}}$ are: $h_{r_{f_{u(k)}}}\left(x_{f_{u(k)}}\right) = h_{r_{f_{u(k)}}}(x_1) + 1$, $h_{r_{f_{u(k)+1}}}\left(x_{f_{u(k)}}\right) = h_{r_{f_{u(k)}}}(x_1)$, and $h_{r_{f_{u(k)}}}\left(x_{f_{u(k)+1}}\right) = h_{r_{f_{u(k)}}}(x_1)$, $h_{r_{f_{u(k)+1}}}\left(x_{f_{u(k)+1}}\right) = h_{r_{f_{u(k)}}}(x_1) + 1$, respectively. Thus, after the move, the congestion level on $r_{f_{u(k)}}$ in $x_{f_{u(k)+1}}$ is the same as in x_{f_v}, while the congestion on $r_{f_{u(k)+1}}$ is greater by 1 than in x_{f_v}. Since $r_{f_{u(k)}} \rightarrow r_{f_{u(k)+1}}$ is a best unit move at $x_{f_{u(k)}}$, reversing this move (i.e., moving back from $x_{f_{u(k)+1}}$ to $r_{f_{u(k)}}$) in $x_{f_{u(k)+1}}$) is not possible. Given this, since the congestion on $r_{f_{u(k)+1}}$ in x_{f_v} is decreased, compared to that in $x_{f_{u(k)+1}}$, the move $x_{f_{u(k)+1}} \rightarrow r_{f_{u(k)}}$ is also not feasible. Note that the proof above also works for the case where $r_{f_{u(k)}} \rightarrow r_{f_{u(k)+1}}$ is the first move of the sequence (although the values of congestion levels are slightly different).

Now let $k \neq l$. Again, we first consider the case where none of the moves $r_{f_{u(k)}} \rightarrow r_{f_{u(k)+1}}$ and $r_{f_{u(l)}} \rightarrow r_{f_{u(l)+1}}$ is the first move of the sequence. If $k < l$ (i.e. the agent makes $r_{f_{u(k)}} \rightarrow r_{f_{u(k)+1}}$ earlier), then consider the move $r_{f_{u(l)}} \rightarrow r_{f_{u(l)+1}}$, and let $x_{f_{u(l)}}$ and $x_{f_{u(l)+1}}$ denote the strategy profiles before and after this move, respectively. Since agent i makes this move later, in $x_{f_{u(l)+1}}$, the congestion level of $r_{f_{u(k)+1}}$ and $r_{f_{u(l)}}$ is the same as they have in x_{f_v}. Given this, the improvement value of move $x_{f_{u(k)+1}} \rightarrow r_{f_{u(l)}}$ is exactly the same as it is in $x_{f_{u(l)+1}}$. Since $r_{f_{u(l)}} \rightarrow r_{f_{u(l)+1}}$ is a best unit move in $x_{f_{u(l)}}$, resource $r_{f_{u(l)}}$ belongs to the set defined in (11) (i.e. set of p^*); that is, reducing a unit from $r_{f_{u(l)}}$ guarantees a maximal cost reduction to agent i in strategy profile $x_{f_{u(l)}}$. This implies that the cost reduction by reducing a unit from $r_{f_{u(k)+1}}$ is not greater than the cost reduction by reducing a unit from $r_{f_{u(l)}}$. Given this, it is easy to see that the reverse move $x_{f_{u(k)+1}} \rightarrow r_{f_{u(l)}}$ in strategy profile $x_{f_{u(l)+1}}$ cannot be positive (i.e., it is not a feasible move). The proof for $k > l$ works in a similar way.

This implies that none of $x_{f_{u(k)+1}} \rightarrow r_{f_{u(l)}}$ is feasible in x_{f_v}. Thus, x_{f_v} satisfies the conditions of the theorem, where x_1 replaces x and x_{f_v} replaces y.

Step 4: Next, consider the case where the sequence of forward moves, $\{r_{f_1}, r_{f_2}, \ldots r_{f_K}\}$, is finite (i.e. $K < \infty$). At this point, we allow agents to make backward moves (i.e., moves that add a unit to r_t^- at each step t). Let $\{r_{b_1}, r_{b_2}, \ldots\}$ denote the sequence of backward moves, where $\forall t$, $r_{b_{t+1}} \rightarrow r_{b_t}$ is the backward move made by some agent i. Note that here $r_{b_1} = r_{f_1}$. Similarly to the case of forward moves, one can show that if there is a loop within $\{r_{b_1}, r_{b_2}, \ldots\}$, then we can leave that loop out from the sequence. Furthermore, one of the following must hold for $\{r_{b_1}, r_{b_2}, \ldots\}$: (i) apart from r_{b_1}, $\{r_{b_1}, r_{b_2}, \ldots\}$ also contains a resource r_{f_u} from the sequence of $\{r_{f_1}, r_{f_2}, \ldots r_{f_K}\}$; that is, $\exists v > 1, u > 0$ such that $r_{b_v} = r_{f_u}$; or (ii) it does not contain such resource, but then it must be finite. The proof is also based on contradiction, and is similar to the proof described in Step 2.

Now, if $\{r_{b_1}, r_{b_2}, \ldots\}$ contains a resource from the sequence of forward moves, then consider the following sequence of moves: $\{r_{f_1}, r_{f_2}, \ldots r_{f_u}, r_{b_1}, r_{b_2}, r_{b_v}\}$. That is, we leave out all the moves after r_{b_v} in the sequence of backward moves, and all the moves after r_{f_u} in the sequence of forward moves. It is easy to see that this sequence is also feasible, that is, all of the moves are best unit moves of some agent i who violates the conditions of the theorem. Now, if the sequence contains inner loops from the backward moves side (the subsequence of forward moves is loopless), we leave these loops out. This way, we obtain a loop similar to the loop described in Step 3. Let x_{b_v} denote the strategy profile resulted by the sequence of moves within $\{r_{f_1}, r_{f_2}, \ldots r_{f_v}, r_{b_1}, r_{b_2}, r_{b_v}\}$. We show that x_{b_v} satisfies the conditions of the theorem (the proof is similar to the one described in Step 3).

Finally, in the case where $\{r_{b_1}, r_{b_2}, \ldots, r_{b_L}\}$ is also finite (i.e. $L < \infty$), let x_T denote the resulting strategy profile by making the moves of the combined sequence $\{r_{f_1}, r_{f_2}, \ldots r_{f_K}\}$ and $\{r_{b_1}, r_{b_2}, \ldots, r_{b_L}\}$. One can easily see that the conditions of the theorem hold for x_T. This completes the proof. □

ISWCG Algorithm. The proof of Theorem 2 suggests a particular dynamic procedure that consists of best single unit moves (Algorithm 1) and arrives at a pure strategy Nash equilibrium from any starting point in finite time. This implies that convex increasing congestion games ISWCG are *weakly–acyclic* (Monderer and Shapley, 1996)—that is, possess an improvement dynamics whose convergence is guaranteed from an arbitrary initial strategy profile.

Theorem 3. *Algorithm 1 finds a pure strategy Nash equilibrium in a given convex increasing ISWCG.*

Proof. The algorithm constructs a sequence of strategy profiles, $\{x_1, x_2, \ldots\}$, such that $\forall t$, x_{t+1} satisfies Theorem 2 with respect to profile x_t (steps $4 - 19$). Then, Corollary 1 implies that $\forall t$, $P(x_t) > P(x_{t+1})$, where $P(x)$ is defined in equation (13). That is, sequence $\{P(x_1), P(x_2), \ldots\}$ is strictly decreasing. Hence, since the game is finite, the algorithm terminates in a PSNE after a finite number of steps. □

5 Conclusions

In this paper, we explore the conditions for PSNE existence in integer–splittalbe congestion games. Although these games do not necessarily admit such an equilibrium, we prove that it is guaranteed to exist in an important subclass of ISWCGs with convex increasing unit–cost functions. Furthermore, we demonstrate that although convex increasing ISWCGs do not have the FIP property, they do possess weak acyclicity, and we provide a natural procedure that achieves an equilibrium from an arbitrary initial strategy profile.

Our results suggest several directions for future research. Specifically, given PSNE existence and convergence, it is important to address further properties of integer–splittable congestion games, such as completeness of the model, quality of solutions

Algorithm 1. ISWCG Algorithm.

1: **Initialisation:** Let $t = 1$, $x_t = x$
2: If $\nexists\, i : D^i_{\max} > 0 \rightarrow$ STOP
3: **while** PSNE not found **do**
4: $x_t \leftarrow$ starting position
5: $\{r_f\} \leftarrow$ sequence of forward moves, $k = 1$
6: **while** forward move is feasible **do**
7: make a forward move: $r_{f_k} \rightarrow r_{f_{k+1}}$, $k := k + 1$
8: **if** there is an inner loop **then** leave out the loop
9: **if** $r_{f_k} = r_{f_1}$ **then** $x_{t+1} \leftarrow$ resulting resource profile of $\{r_{f_1}, r_{f_2}, \ldots r_{f_k}\}$ from x_t, GOTO STEP 20
10: **end while**
11: $\{r_{f_1}, r_{f_2}, \ldots r_{f_K}\} \leftarrow$ resulting sequence of forward moves
12: $\{r_b\} \leftarrow$ sequence of backward moves, $l = 1$
13: **while** backward move is feasible **do**
14: make a backward move: $r_{b_{l+1}} \rightarrow r_{b_l}$, $l := l + 1$
15: **if** there is an inner loop **then** leave out the loop
16: **if** $\exists r_{f_v} \in \{r_f\}$ such that $r_{b_l} = r_{f_v}$ **then** $x_{t+1} \leftarrow$ resulting resource profile of $\{r_{f_1}, r_{f_2}, \ldots r_{f_u}, r_{b_1}, r_{b_2}, r_{b_v}\}$ from x_t, GOTO STEP 20
17: **end while**
18: $\{r_{b_1}, r_{b_2}, \ldots, r_{b_L}\} \leftarrow$ resulting sequence of backward moves
19: $x_{t+1} \leftarrow$ resulting resource profile of $\{r_{f_1}, r_{f_2}, \ldots r_{f_K}\}$ and $\{r_{b_1}, r_{b_2}, \ldots, r_{b_L}\}$ from x_t

20: **if** $x_{t+1} =$ PSNE **then** STOP
21: $t := t + 1$
22: **end while**

and computational complexity. To this end, we aim to (i) investigate how far the assumptions on the agents' strategy sets and payoff functions can be relaxed while still guaranteeing the existence of pure strategy equilibria, (ii) characterise the efficiency of PSNE in terms of prices of anarchy and stability, and (iii) provide a complexity analysis of the problem of finding equilibria and develop efficient algorithmic solutions, if applicable.

References

Beier, R., Czumaj, A., Krysta, P., Vöcking, B.: Computing equilibria for congestion games with (im)perfect information. In: Proc. of SODA 2004, pp. 746–755 (2004)

Byde, A., Polukarov, M., Jennings, N.R.: Games with congestion-averse utilities. In: Mavronicolas, M., Papadopoulou, V.G. (eds.) SAGT 2009. LNCS, vol. 5814, pp. 220–232. Springer, Heidelberg (2009)

Cominetti, R., Correa, J.R., Stier-Moses, N.E.: The impact of oligopolistic competition in networks. Operations Research 57(6), 1421–1437 (2009)

Dunkel, J., Schulz, A.: On the complexity of pure-strategy nash equilibria in congestion and local-effect games. Mathemathics of Operations Research 33(4), 851–868 (2008)

Koutsoupias, E., Papadimitriou, C.H.: Worst–case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, p. 404. Springer, Heidelberg (1999)

Krysta, P., Sanders, P., Vöcking, B.: Scheduling and traffic allocation for tasks with bounded splittability. In: Proc. of the 28th International Symposium on Mathematical Foundations of Computer Science, pp. 500–510 (2003)

Leslie, D.S., Collins, E.J.: Generalised weakened fictitious play. Games and Economic Behavior 56, 285–298 (2006)

Leyton-Brown, K., Tennenholtz, M.: Local-effect games. In: Proc. of IJCAI 2003, pp. 772–780 (2003)

Marden, J.R., Arslan, G., Shamma, J.S.: Regret based dynamics: Convergence in weakly acyclic games. In: Proc. of AAMAS 2007, pp. 194–201 (2007)

Meyers, C.: Network flow problems and congestion games: complexity and approximation results. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, USA. AAI0809430 (2006)

Milchtaich, I.: Congestion games with player–specific payoff functions. Games and Economic Behavior 13(1), 111–124 (1996)

Monderer, D., Shapley, L.S.: Potential games. Games and Economic Behavior 14, 124–143 (1996)

Orda, A., Rom, R., Shimkin, N.: Competitive routing in multiuser communication networks. IEEE/ACM Trans. Networking 1(5), 510–521 (1993)

Penn, M., Polukarov, M., Tennenholtz, M.: Congestion games with load-dependent failures: Identical resources. Games and Economic Behavior 67(1), 156–173 (2009a)

Penn, M., Polukarov, M., Tennenholtz, M.: Random order congestion games. Mathematics of Op. Res. 34(3), 706–725 (2009b)

Rosen, J.: Existence and uniqueness of equilibrium points for concave n-person games. Econometrica 33, 520–534 (1965)

Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. Int. J. of Game Theory 2, 65–67 (1973a)

Rosenthal, R.W.: The network equilibrium problem in integers. Networks 3, 53–59 (1973b)

Roughgarden, T., Tardos, E.: How bad is selfish routing? J. ACM 49(2), 236–259 (2002)

Shachnai, H., Tamir, T.: Multiprocessor scheduling with machine allotment and parallelism constraints. Algorithmica 32, 651–678 (2002)

Voice, T.D., Polukarov, M., Jennings, N.R.: On the impact of strategy and utility structures on congestion-averse games. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 600–607. Springer, Heidelberg (2009)