
Equilibria Computation in Multidimensional Congestion Games: CSP and Learning Dynamics Approaches

Mohammad T. Irfan¹

Hau Chan²

Jared Soundy²

¹Department of Computer Science, Bowdoin College, Brunswick, Maine, USA

²School of Computing, University of Nebraska-Lincoln, Lincoln, Nebraska, USA

Abstract

We present algorithms of two flavors—one rooted in constraint satisfaction problems (CSPs) and the other in learning dynamics—to compute pure-strategy Nash equilibrium (PSNE) in k -dimensional congestion games (k -DCGs) and their variants. We first show that deciding the existence of a PSNE in a k -DCG is NP-complete even when players have binary and unit demand vectors. For general cost functions (potentially non-monotonic), we devise a new CSP-inspired algorithmic framework for PSNE computation, leading to algorithms that run in polynomial time under certain assumptions while offering exponential savings over standard CSP algorithms. We further refine these algorithms for variants of k -DCGs. Our experiments demonstrate the effectiveness of this new CSP framework for hard, non-monotonic k -DCGs. We then provide learning dynamics-based PSNE computation algorithms for linear and exponential cost functions. These algorithms run in polynomial time under certain assumptions. For general cost, we give a learning dynamics algorithm for an (α, β) -approximate PSNE (for certain α and β). Lastly, we also devise polynomial-time algorithms for structured demands and cost functions.

1 INTRODUCTION

In non-cooperative games, a player’s payoff depends on their own choice of action and the choices of actions by the other players. In general, the payoff may change depending on *who* chose a particular action. In a seminal paper, Rosenthal presented a special class of games—to become famously known as *congestion games* later—where the number of players rather than the identities of the players choosing an action is relevant [Rosenthal, 1973]. In a congestion game,

there is a set of resources (e.g., edges in a road network). Each player has a set of strategies, where each strategy is a subset of resources (e.g., paths in a network). A strategy profile consists of a strategy for each player. The cost of a resource (e.g., edge) is a function of the number of players using that resource. Given a strategy profile, a player’s cost is the sum of the costs of the resources used by the player. A strategy profile is a pure-strategy Nash equilibrium (PSNE) if no player has incentive to deviate unilaterally.

The congestion games literature can be divided into three main frontiers: unweighted, weighted one-dimensional, and weighted multidimensional. Unweighted congestion games are the classical ones [Rosenthal, 1973], where the guaranteed existence of a PSNE naturally leads to computational questions. In their seminal work, Monderer and Shapley [1996] showed that any unweighted congestion game is a *potential game*, which is appealing for learning dynamics. For unweighted congestion games on networks, if the game is *symmetric* (same start-end pair for all), then there exists a polynomial-time network-flow algorithm to find a PSNE; otherwise, the problem is PLS-complete [Fabrikant et al., 2004], even for linear cost [Ackermann et al., 2008].

In a weighted congestion game, each player has a weight or demand, and the cost of a resource is a function of the sum of the demands of the players using that resource. Unlike unweighted congestion games, a PSNE is not guaranteed to exist in weighted congestion games [Libman and Orda, 1997, Fotakis et al., 2005]. Dunkel and Schulz [2008] went one step further and showed that PSNE existence in weighted congestion games is strongly NP-complete, even for a constant number of players.

On the positive side, a PSNE is guaranteed to exist in a weighted congestion game when the cost function is linear [Fotakis et al., 2005] or exponential [Harks and Klimm, 2012]. Harks et al. [2011] characterized the existence of potential functions. Special cases involving parallel edges have also received attention [Milchtaich, 1996, Fotakis et al., 2002, Gairing et al., 2004, Mavronicolas et al., 2007].

Table 1: Our main results on k -dimensional congestion games (k -DCGs), k -class congestion games (k -CCGs), and variants. Notation: NPC \equiv NP-Complete, $n = \#$ players, $m = \#$ resources, $p = \max \#$ strategies, $\mathbf{d}_i =$ player i 's demand vector, $\mathbf{d}_N = \sum_i \mathbf{d}_i$, $w_{\max} = \max_j \mathbf{d}_{N_j}$, $\tilde{n} = \max \#$ players selecting a resource in a binary k -DCG, or $\max \#$ players of a type in a k -DCG with player types, $l(i) =$ nonzero-element index in \mathbf{d}_i for k -CCG, a_{\max} , b_{\max} , and \mathbf{z} are cost parameters.

† We give approximation algorithms for (α, β) -PSNE, which always exists. ‡ Klimm and Schütz [2022].

	Problem	PSNE	Time Complexity to Determine or Compute PSNE
CSP Framework	General Cost k -DCG	NPC†	$\mathcal{O}((w_{\max})^{km}(nkp^2m^2 + nkmp(w_{\max})^{km}))$
	Subclass: Binary k -DCG	NPC	$\mathcal{O}(\tilde{n}^{km}(nkp^2m^2 + \min\{nkmp\tilde{n}^{km}, n^{km+1}p\}))$
	Subclass: k -CCG	NPC	$\mathcal{O}((w_{\max})^{km}(np^2m^2 + nkpm(w_{\max})^m))$
	Subclass: k -DCG with player types	NPC	$\mathcal{O}((\tilde{n})^{\tau m}(np^2m^2 + n\tau pm(\tilde{n})^m) + \tau nk)$
Learning Dynamics	Linear Cost k -DCG	Always‡	$\mathcal{O}\left(nkpm^2 \times n^2m(a_{\max} + b_{\max}) \frac{\max_i [\mathbf{z} \cdot \mathbf{d}_i]^2}{\min_i [\mathbf{z} \cdot \mathbf{d}_i]}\right)$
	Linear Subclass: Binary k -DCG	Always	$\mathcal{O}\left(nkpm^2 \times n^2m(a_{\max} + b_{\max})(k \max_j z_j)^2\right)$
	Linear Subclass: k -CCG	Always	$\mathcal{O}\left(nkpm^2 \times n^2m(a_{\max} + b_{\max}) \frac{\max_j z_j^2}{\min_j z_j} \frac{\max_i d_{i,l(i)}^2}{\min_i d_{i,l(i)}}\right)$
	Exponential Cost k -DCG	Always‡	$\mathcal{O}\left(nkpm^2 \times \frac{e}{e-1} (m \exp(\mathbf{z} \cdot \mathbf{d}_N) a_{\max} + nmb_{\max})\right)$
Structured	Ordered \mathbf{d}_i 's, nondec. cost, singleton str.	Always	$\mathcal{O}(n \log n + nmk)$
	Ordered \mathbf{d}_i 's, nondec. cost, shared str.	Always	$\mathcal{O}(n \log n + npmk)$
	Structured cost, singleton str.	Always	$\mathcal{O}(n \log n + nmk)$

Multidimensional congestion games is a very recent frontier which we investigate here. Introduced by Klimm and Schütz [2014], this class of games is a generalization of weighted congestion games where the demand of each player is a k -dimensional vector. Very recently, Klimm and Schütz [2022] have shown that certain affine and exponential cost functions are the only ones for which a PSNE exists for sure. Their characterization leads to the following *computational* questions investigated here: *How can we compute a PSNE (if it exists) in multidimensional congestion games and their variants? How hard is this computation?*

These questions are motivated by many real-world applications. Advances in multidimensional congestion games may contribute to richer traffic models that account for the heterogeneity in vehicles (e.g., weight, length, etc.). Such multidimensional models were envisioned by transportation researchers many decades ago [Dafermos, 1972] and are now topics of active investigation [Van Lint et al., 2008, Pi et al., 2019, Wang et al., 2019]. Computational advances may also contribute to various other application areas—wireless networks [Yamamoto, 2015], distributed systems [Nadig et al., 2022, 2019], telecommunication [Altman et al., 2006], and smart grids [Fadlullah et al., 2011], to name just a few.

Our Contributions

Driven by whether or not a PSNE is guaranteed to exist, we take two fundamentally different computational approaches inspired by CSPs and learning dynamics. The CSP approach can handle *any* k -DCGs (for which a PSNE may not exist), whereas learning dynamics can handle certain k -DCGs with a PSNE. Table 1 summarizes our main results.

For general k -DCGs, we devise a CSP whose dual decouples the players' strategies. We give algorithms that utilize this decoupling and run in polynomial time under certain assumptions (Section 4). *To our knowledge, this CSP framework is new within the rich congestion games literature spanning over five decades.* The significance of our CSP framework lies not only in the exponential savings it offers compared to well-known CSP algorithms but also in its applicability beyond congestion games.

For linear and exponential cost, we give iterative learning dynamics algorithms for k -DCGs and their variants by deriving and bounding weighted potential functions (Section 5). For general cost, we show that for certain α and β , there is always an (α, β) -approximate PSNE that can be computed via learning dynamics. We also give polynomial-time algorithms for structured cost and demands (Section 6).

The significance of our computational results can be best understood against the backdrop of hardness results (Section 3). We show that deciding the existence of a PSNE in a k -DCG is NP-complete for very special cases. Put together, this paper addresses computational questions while giving new insights for some extremely hard problems.

2 PRELIMINARIES

We formally define multi-dimensional congestion games and related game-theoretic terms. Roughly speaking, a multi-dimensional congestion game is a natural generalization of weighted congestion games where the weight or demand of each player is a multidimensional vector. The cost of

each resource is a function of the aggregated demands of the players using that resource.

More formally, a k -dimensional congestion game (k -DCG) consists of a set $N = \{1, \dots, n\}$ of n players and a set $R = \{1, \dots, m\}$ of m resources. Each player $i \in N$ has two elements: (1) a strategy set $S_i \subseteq 2^R \setminus \{\emptyset\}$, defined to be subsets of resources that i can select and (2) a k -dimensional demand vector $\mathbf{d}_i = (d_{i_1}, \dots, d_{i_k}) \in \mathbb{R}^k$, consisting of the weight or demand of player i at each dimension $1, \dots, k$. Each resource $r \in R$ has a cost function $c_r : \mathbb{R}^k \rightarrow \mathbb{R}$ that maps k -dimensional real-valued vectors to real numbers. We use $p = \max_{i \in N} |S_i|$ to denote the maximum number of strategies for any player.

Given a strategy profile $\mathbf{s} = (s_1, \dots, s_n) \in S = S_1 \times \dots \times S_n$ of n players, let $\mathbf{x}_r(\mathbf{s}) = \sum_{i \in N: r \in s_i} \mathbf{d}_i$ be the aggregated k -dimensional demand vector of the players who select resource r under the strategy profile $\mathbf{s} \in S$. Naturally, given a strategy profile \mathbf{s} , the cost function of player i is defined to be $\pi_i(\mathbf{s}) = \pi_i(s_i, \mathbf{s}_{-i}) = \sum_{r \in s_i} c_r(\mathbf{x}_r(\mathbf{s}))$, i.e., the sum of the costs of the resources selected by player i under s_i , given others' strategies \mathbf{s}_{-i} .

We are interested in computing PSNE in k -DCGs and their variants listed below. We present these variants with motivating examples from the domain of load balancing in distributed systems [Nadig et al., 2022, 2019, Anantha et al., 2017]. k -DCGs naturally model various dimensions of user demands in distributed systems, such as bit rates, latency, error tolerance, and throughputs.

- k -DCGs with binary demand vectors $\mathbf{d}_i \in \{0, 1\}^k \forall i$. Example: data flow in distributed systems can be short-lived or long-lived, bursty or deterministic, etc.
- k -class congestion games (k -CCGs), where each demand vector has one positive element, the rest being zeros. Example: different use-cases, such as streaming, video conferencing, web browsing, etc.
- k -DCGs with player types, where players of the same type are characterized by the same demand vector. Example: categories of traffic on a campus network: VPN, student access, scientific computation, etc.

We next define PSNE and approximate PSNE—two solution concepts of our interest.

Definition 1. (*Pure-Strategy Nash Equilibrium (PSNE)*) A strategy profile $\mathbf{s}^* = (s_1^*, \dots, s_n^*) \in S$ is a pure-strategy Nash equilibrium (PSNE) in a k -DCG if and only if for each player $i \in N$ and any $s'_i \in S_i$, we have that $\pi_i(\mathbf{s}^*) \leq \pi_i(s'_i, \mathbf{s}_{-i}^*)$.

Definition 2. (*(α, β) -PSNE*) A strategy profile $\mathbf{s} = (s_1, \dots, s_n) \in S$ is an (α, β) -approximate PSNE in a k -DCG for some $\alpha \geq 1$ and $\beta \geq 0$ if and only if for each player $i \in N$ and any $s'_i \in S_i$, we have that

$\pi_i(\mathbf{s}) \leq \alpha \pi_i(s'_i, \mathbf{s}_{-i}) + \beta$. When we mention α -PSNE (without β), we mean $\beta = 0$.

Constraint Satisfaction Problem (CSP)

A CSP is specified by a set of *variables*, a *domain* for each variable, and a set of *constraints*, each constraint being over a subset of variables known as its *scope*. A CSP asks us to assign a value to each variable from their respective domains so that all the constraints are satisfied. A wide range of problems, such as Boolean satisfiability, map coloring, scheduling, and even PSNE computation in games, can be modeled as CSPs [Dechter, 2003, Gottlob et al., 2003].

We often represent the structural information of a (primal) CSP using a *primal constraint network*, where each node represents a variable, and each edge connects two variables that appear together in a constraint (potentially with other variables). As a result, unless the constraints are binary, we cannot identify the scope of a constraint just by looking at the primal constraint network.

A CSP also has a *dual constraint network*, where each variable represents a constraint, and each edge connects two constraints with shared variables in their scopes and is labeled with these shared variables. The dual constraint network leads to the *dual CSP*, where the domain of each dual variable is computed as follows: Consider its corresponding primal constraint and assign values to the scope of the primal constraint to satisfy it. Such assignments constitute the domain of the dual variable. Furthermore, the dual CSP enforces the edge-wise dual constraint that each primal variable shared between any two dual variables must have the same value in both. Therefore, the dual CSP is a reformulation of the primal CSP and contains only binary constraints.

3 COMPUTATIONAL COMPLEXITY

We show that deciding the existence of a PSNE in special variants of k -DCGs is NP-complete. The NP-hardness of general k -DCGs is not surprising because determining a PSNE in weighted congestion games (i.e., when $k = 1$) is already strongly NP-complete [Dunkel and Schulz, 2008].

What is surprising is that we show that determining the existence of a PSNE in k -DCGs is NP-complete even when each player i 's k -dimensional demand vector \mathbf{d}_i is a binary vector (even a unit vector) for some polynomially bounded k . In sharp contrast, there is always a PSNE in unweighted (1-dimensional) congestion games [Rosenthal, 1973]. Furthermore, if the players have the same demand vector, the game is guaranteed to have a PSNE by reducing it to an unweighted congestion game. We have the following result.

Theorem 3. *Deciding the existence of a PSNE in a k -DCG is NP-complete even when the demand vector \mathbf{d}_i of each player $i \in N$ is a binary vector and k is sublinear in the number of players. That is, $\mathbf{d}_i \in \{0, 1\}^k$ for all i and $k =$*

$\mathcal{O}(\log n)$.

Proof Sketch. The problem is in NP because verifying that a strategy profile $\mathbf{s}^* \in S$ is a PSNE takes polynomial time. For NP-hardness, we reduce from weighted congestion games [Dunkel and Schulz, 2008]. Given a weighted congestion game we construct a k -DCG with identical sets of players, resources, and actions. In the k -DCG game we give the players binary demand vectors equivalent to the binary representations of the integer weights from the weighted congestion game. The length of the demand vector is set to $k = \lfloor \log \max_{i \in N} \tilde{d}_i \rfloor + 1$ where \tilde{d}_i is the integer weight of player i in the weighted congestion game. Finally, we construct cost functions for k -DCG that we show to yield the same cost given the same strategy profile for all agents. Therefore, a strategy profile is a PSNE in one game if and only if it is a PSNE in the other game. \square

Next, we investigate whether PSNE computation is easier for restricted demands. Unfortunately, even when the binary demand vector is a unit vector, the problem remains hard.

Theorem 4. *Deciding the existence of a PSNE in a k -DCG (or a k -CCG) is NP-complete even when the demand vector \mathbf{d}_i of each player $i \in N$ is a binary unit vector and k is linear of the number of players. That is, $\mathbf{d}_i \in \{\mathbf{x} \in \{0, 1\}^k; \sum_{j=1}^k x_j = 1\}$ for all i and $k = \mathcal{O}(n)$.*

Proof Sketch. The problem is clearly in NP. The NP-hardness reduction is from weighted congestion games. \square

4 GENERAL COST: A CSP APPROACH

We can formulate the PSNE computation problem in a k -DCG as a CSP, which consists of (1) a variable for each player, (2) the domain of a variable being the corresponding player's strategy set, and (3) a best-response constraint for each player i , representing i 's best responses s_i to any \mathbf{s}_{-i} .

As illustrated in Fig. 1 (a) and (b), the nature of the n -ary best-response constraints means that both the primal and the dual constraint networks are complete networks. Furthermore, *all* players appear on each edge of the dual network. This portrays a grim picture because it is hard to design efficient algorithms without decoupling the players' strategies. For example, one solution approach is to check each strategy profile for a PSNE by verifying Definition 1. Letting $p = \max_{i \in N} |S_i|$, this approach takes $\mathcal{O}(np^{n+1})$ time, which is exponential in the number of players.

The grave computational implication of not decoupling the players' strategies leads us to a key technical insight. Instead of using the above CSP, we first construct a different CSP for k -DCGs and then consider its dual. In the new CSP, the variables are the players and the *configuration* Y of the game.

The domain of each player i is their strategy set S_i and that of Y is the set of all k -dimensional aggregated demand vectors for m resources, $\mathbf{y} \equiv (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$. There are n binary constraints, each capturing a player's *best response to a configuration*. We use the structure of k -DCGs to define such best responses: For any configuration \mathbf{y} , a player i 's best-response strategies are $s_i \in S_i$ that minimize the cost $\sum_{r \in s_i} c_r(\mathbf{y})$. There is an additional *feasibility constraint* that enforces that the strategy profile \mathbf{s} assigned to the players are consistent with the aggregated demand vectors \mathbf{y} assigned to Y ; i.e., $\mathbf{x}_1(\mathbf{s}) = \mathbf{y}_1, \mathbf{x}_2(\mathbf{s}) = \mathbf{y}_2, \dots, \mathbf{x}_m(\mathbf{s}) = \mathbf{y}_m$. The primal constraint network for this CSP is shown in Fig. 1(c) and its dual in Fig. 1(d). Most notably, the dual CSP allows us to decouple the players' strategies from each other.

To our knowledge, this dual CSP, which grounds our algorithmic framework, has not been studied in the congestion games literature spanning over five decades. To formalize this dual CSP, each dual node is a primal constraint. So, there is a dual node $v_{i,Y}$ for each player i 's best response to the configuration variable Y , and there is one dual variable $v_{N,Y}$ for the primal feasibility constraint (see Fig. 1(d)). For each $i \in N$, there is an edge between $v_{N,Y}$ and $v_{i,Y}$ labeled i, Y . For any $i \neq j \in N$, there is an edge between $v_{i,Y}$ and $v_{j,Y}$ labeled Y . As described in Section 2, each dual variable has a domain consisting of satisfying assignments for the corresponding primal constraint, and the edges in the dual constraint network lead to dual constraints that ensure that the shared primal variables across any edge are assigned the same value in both endpoints of the edge.

Before presenting our algorithmic framework, we show that the dual CSP has a solution if and only if there is a PSNE. To see why, note that the assignments (s_i, \mathbf{y}) made to the $v_{i,Y}$ variables capture the players' best responses to \mathbf{y} , and the edge label between any two $v_{i,Y}$ and $v_{j,Y}$ variables enforces sharing the same \mathbf{y} in these assignments. Furthermore, the assignment $(\mathbf{s}, \mathbf{y}')$ made to the $v_{N,Y}$ variable makes sure that the strategy profile \mathbf{s} leads to the configuration \mathbf{y}' , and the labels on the edges connecting $v_{N,Y}$ to $v_{i,Y}$ enforce that $\mathbf{s} = (s_1, \dots, s_n)$ and $\mathbf{y} = \mathbf{y}'$.

Our algorithmic framework consists of two procedures. Procedure 1 computes the domains of each $v_{i,Y}$ dual variable and Procedure 2 searches for a solution using the computed domains. As a preview, our algorithms are polynomial in n (the number of players), p (the maximum number of strategies for any player), and a maximum weight term when k (number of dimensions) and m (number of resources) are bounded. This is useful when the number of resources and strategies is constant but the number of players can be large. In fact, even with a constant number of players, determining PSNE existence a weighted congestion game is already strongly NP-complete [Dunkel and Schulz, 2008].

As Fig 1(d) shows, Y is shared across all edges. Therefore,




Figure 1: Our key technical insight: **(a)** Typical CSP for $N = \{1, \dots, 4\}$: Each node is a player with their strategy set as the domain, and each player has a best-response constraint involving all the players. **(b)** The dual of (a): Each dual node BR_i represents player i 's best-response constraint with its domain being strategy profiles (s_i, \mathbf{s}_{-i}) where s_i is i 's best response to \mathbf{s}_{-i} . Each edge shows that *all* players are shared between its endpoints, which makes it hard to decouple the strategies of the players. **(c)** A new CSP we present where in addition to the players, there is a node Y for the configuration. The constraints are: (1) each player i plays its best response to Y and (2) the strategies assigned to the players are consistent with the configuration assigned to Y . **(d)** The dual of (c): Edges show decoupling of the players' strategies. Contrast it with (b).

we parameterize our algorithms by any configuration given as input. This leads to the question of how many configurations there can be. The demand vector $\mathbf{d}_i = (d_{i_1}, \dots, d_{i_k})$ of each player i being an integer vector (standard assumption [Dunkel and Schulz, 2008]), we define $w_j = \sum_{i \in N} d_{i_j}$ for each $j = 1, \dots, k$. Letting $w_{\max} = \max_{j \in [k]} w_j$, we have $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m \in \{0, \dots, w_{\max}\}^k$. Thus, we only need to consider at most $(w_{\max} + 1)^{km}$ or $\mathcal{O}((w_{\max})^{km})$ configurations. We are now ready for the algorithms. Please see the Appendix for pseudocode.

Procedure 1: Compute Domains of Dual Variables $v_{i,Y}$
 Given a configuration $\mathbf{y} \equiv (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$, where each $\mathbf{y}_j \in \{0, \dots, w_{\max}\}^k$, we compute the set of strategies for each player i that makes i "happy" under the configuration. To do this, abusing the notation π_i slightly, we define and compute, for any $i \in N$, $s_i, s'_i \in S_i$, and $s_i \neq s'_i$,

$$\begin{aligned} \pi_i(s_i, \mathbf{y}) &= \sum_{r \in s_i} c_r(\mathbf{y}_r) \\ \pi_i(s_i, \mathbf{y}, s'_i) &= \sum_{r \in s'_i \cap s_i} c_r(\mathbf{y}_r) + \sum_{r \in s'_i \setminus s_i} c_r(\mathbf{y}_r + \mathbf{d}_i) \\ BR_i(\mathbf{y}) &= \{s_i \in S_i \mid \forall s'_i \in S_i, \pi_i(s_i, \mathbf{y}) \leq \pi_i(s_i, \mathbf{y}, s'_i)\} \end{aligned}$$

The first equation calculates player i 's cost. The second calculates player i 's cost when deviating from s_i (under \mathbf{y}) to s'_i . $BR_i(\mathbf{y})$ computed in the last equation is the set of i 's best responses to \mathbf{y} . Therefore, the domain of $v_{i,Y}$ is the union of sets $\{(s_i, \mathbf{y}) \mid s_i \in BR_i(\mathbf{y})\}$ for all \mathbf{y} .

We deliberately do not compute the domain of $v_{N,Y}$ (the dual variable for the primal feasibility constraint) because it may contain numerous strategy profiles that are not PSNE. We next show how we can search for PSNE without explicitly computing the domain of $v_{N,Y}$.

Procedure 2: Search for PSNE

Given a configuration $\mathbf{y} \equiv (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$, a PSNE under it is a strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ such that

(1) (s_i, \mathbf{y}) is in the domain of $v_{i,Y}$ for each player i , and
 (2) $\mathbf{x}_1(\mathbf{s}) = \mathbf{y}_1, \mathbf{x}_2(\mathbf{s}) = \mathbf{y}_2, \dots, \mathbf{x}_m(\mathbf{s}) = \mathbf{y}_m$. The first condition enforces players' best responses to \mathbf{y} , while the second condition enforces the primal feasibility constraint. We present the following general result.

Theorem 5. *For any k -DCG, there is an algorithm to determine the existence of a PSNE in $\mathcal{O}((w_{\max})^{km}(nkp^2m^2 + nkm p(w_{\max})^{km}))$. The algorithm is polynomial in n , p , and w_{\max} , when m and k are constants.*

Proof Sketch. Procedure 1 runs in $\mathcal{O}(nkp^2m^2)$ for all players. Procedure 2 can be done efficiently using dynamic programming (DP), where we (1) first order the players $1, \dots, n$ and (2) create a binary table $T_i(\mathbf{y}'_1, \mathbf{y}'_2, \dots, \mathbf{y}'_m) \in \{0, 1\}$ for each $\mathbf{y}'_1, \mathbf{y}'_2, \dots, \mathbf{y}'_m \in \{0, \dots, w_{\max}\}^k$ of size $\mathcal{O}((w_{\max})^{km})$ for each player i . We first initialize $T_0(\mathbf{0}, \dots, \mathbf{0}) = 1$ where we have an all zero configuration. We then define $T_i(\mathbf{y}'_1, \mathbf{y}'_2, \dots, \mathbf{y}'_m) = 1$ if and only if there is $\bar{\mathbf{y}}_1, \bar{\mathbf{y}}_2, \dots, \bar{\mathbf{y}}_m$ such that $T_{i-1}(\bar{\mathbf{y}}_1, \bar{\mathbf{y}}_2, \dots, \bar{\mathbf{y}}_m) = 1$ and for some $s_i \in BR_i(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$, $\mathbf{y}'_r = \bar{\mathbf{y}}_r + \mathbb{1}[r \in s_i] \mathbf{d}_i$ for each $r \in R$. Table T_i can be constructed by looking at all the 1 entries of T_{i-1} and adding the player demand vector to the corresponding resources for each $s_i \in BR_i(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$. The DP runs in $\mathcal{O}(nkm p(w_{\max})^{km})$. \square

To see the significance of the above result, note that we can use well-known algorithms to solve the dual CSP. For instance, backtracking algorithms with graph-based learning can solve the dual CSP in $\mathcal{O}\left((n+1)^2 \cdot (2 \cdot p^n \cdot w_{\max}^{km})^{n+1}\right)$ time, which is exponential in n [Dechter, 2003][Ch 6]. In contrast, *our algorithm guarantees an exponential saving.*

Variants: Binary Demand Vectors

In Section 3, we showed that k -DCGs with binary demand vectors are provably hard. We can still apply Theorem 5 to derive a pseudopolynomial time algorithm when k and m are bounded. However, an improved analysis gives us

the following result. Note that in the case of binary demand vectors, any j -th element of an aggregated demand vector corresponds to the number of players having the j -th bit of their demand vector “on.” Therefore, for clarity, we use $\tilde{n} = \max_{j \in [k]} \sum_{i \in N} d_{i,j}$ in place of w_{\max} to denote the maximum number of players having a demand vector bit on. The following result is particularly interesting when $\tilde{n} \ll n$.

Theorem 6. *For k -DCGs with binary demand, there is an $\mathcal{O}(\tilde{n}^{km}(nkp^2m^2 + \min\{nkpmp\tilde{n}^{km}, n^{km+1}p\}))$ -time algorithm to compute a PSNE or decide none exists. The algorithm is polynomial in n and p when m and k are constants.*

Proof Sketch. Putting $w_{\max} = \tilde{n}$ in Theorem 5, the running time is $\mathcal{O}(\tilde{n}^{km}(nkp^2m^2 + nkpmp\tilde{n}^{km}))$. However, using a different analysis that exploits the bit-vector structure, we can shave off a factor of km from the second term at the expense of having n^{km} instead of \tilde{n}^{km} . This would be useful when $\tilde{n} \approx n$. The main idea is when we consider player i in Procedure 2, the number of configurations for T_i is at most $(i+1)^{km}$, leading to $\mathcal{O}(\sum_{i=1}^n [(i+1)^{km} + kpmi^{km}])$ or $\mathcal{O}(n^{km+1}p)$ time for Procedure 2. \square

Variant: k -Class Congestion Game (k -CCG)

Let the class of player i be the index where the positive element appears in \mathbf{d}_i . Although Theorem 5 can be directly applied to this case, we can exploit the structure of the game to improve the running time. The key intuition is that the players can be partitioned according to their classes. The players in a class $j \in [k]$ can only affect the j -th index of the aggregated demand on any resource. That is, they affect the j -th index of each of $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m$. As a result, Procedure 2 can be broken into k different computational tasks, each corresponding to a class. This idea leads us to the following result. *Notably, compared to Theorem 5, this partition-based algorithm removes a k term from the exponent.*

Theorem 7. *For k -CCGs, there is an $\mathcal{O}((w_{\max})^{km}(np^2m^2 + nkpmp(w_{\max})^m))$ algorithm to compute a PSNE or decide none exists. The algorithm is polynomial in n , p , and w_{\max} when m and k are constants.*

Proof Sketch. As a preprocessing step, we partition the players into C_1, \dots, C_k based on their classes. We now do the following operations in each partition C_j independently. We start the DP by ordering the players in C_j as $1, 2, \dots, |C_j|$ (wlog). We then create a binary table $T_i(z_1, z_2, \dots, z_m) \in \{0, 1\}$ for each $z_1, z_2, \dots, z_m \in \{0, \dots, w_{\max}\}$ of size $\mathcal{O}((w_{\max})^m)$ for each player i in C_j . We initialize $T_0(0, \dots, 0) = 1$. We then define $T_i(z_1, z_2, \dots, z_m) = 1$ if and only if there is z'_1, z'_2, \dots, z'_m such that $T_{i-1}(z'_1, z'_2, \dots, z'_m) = 1$ and for some $s_i \in BR_i(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$, $z_r = z'_r + \mathbb{1}[r \in s_i]d_{i,j}$ for each $r \in R$. We have a PSNE if and only if for each partition C_j , $T_{|C_j|}(y_{1j}, y_{2j}, \dots, y_{mj}) = 1$. \square

We next consider the special case of k -CCGs with binary demand vectors (i.e., exactly one bit is “on” in each player’s demand vector). This will be useful when we consider player types next. We get the following corollary from Theorems 6 and 7. Once again, the result is interesting when $\tilde{n} \ll n$.

Corollary 8. *For k -CCGs with binary demand vectors, there is an $\mathcal{O}((\tilde{n})^{km}(np^2m^2 + nkpmp(\tilde{n})^m))$ -time algorithm to compute a PSNE or decide none exists. The algorithm is polynomial in n and p when m and k are constants.*

Variant: k -DCG with Player Types

To motivate this variant, consider a road-traffic setting. There are different types of vehicles, and vehicles of the same type share similarities in their demand vectors. We define players to be of the same type if their demand vectors are the same. Although this setting is very natural, to our knowledge, it has not been fully explored in the literature. Here, other than player types, we do not make any assumptions about the demands or cost functions. While this variant is NP-hard (reduction from k -DCG by making a type for each player), the following result is very appealing when the maximum number of players of any type $\tilde{n} \ll n$.

Theorem 9. *Given a k -DCG with τ types of players and at most \tilde{n} players of any type, there is an $\mathcal{O}((\tilde{n})^{\tau m}(np^2m^2 + n\tau pm(\tilde{n})^m) + \tau nk)$ time algorithm to compute a PSNE or decide that there exists none. The algorithm is polynomial in n and p for bounded m and τ .*

Proof. Let $(N, R, \{S_i, \mathbf{d}_i\}_{i \in N}, \{c_r\}_{r \in R}, k)$ be a k -DCG instance with τ types of players. We reduce this instance to a PSNE-equivalent τ -DCG instance $(N, R, \{S_i, \tilde{\mathbf{d}}_i\}_{i \in N}, \{\tilde{c}_r\}_{r \in R}, \tau)$ as follows. First, we partition the k -DCG players into τ types and store the k -dimensional demand vector (from k -DCG) of any player of type t into $\bar{\mathbf{d}}_t$ (i.e., $\bar{\mathbf{d}}_t = \mathbf{d}_i$ if player i is of type t). This takes $\mathcal{O}(\tau nk)$ time. For each player i of type t in τ -DCG, we define a τ -dimensional unit demand vector $\tilde{\mathbf{d}}_i$ where only the t -th element is 1, the rest being 0. Given a τ -dimensional aggregated demand vector $\tilde{\mathbf{x}}_r(\mathbf{s}) = (\tilde{\mathbf{x}}_r(\mathbf{s})_1, \dots, \tilde{\mathbf{x}}_r(\mathbf{s})_\tau)$, where any t -th element represents the total number of players of type t using r , we define the cost function $\tilde{c}_r(\tilde{\mathbf{x}}_r(\mathbf{s})) = c_r(\sum_{t=1}^{\tau} (\tilde{\mathbf{x}}_r(\mathbf{s}))_t \bar{\mathbf{d}}_t)$. Thus, $\tilde{c}_r(\tilde{\mathbf{x}}_r(\mathbf{s})) = c_r(\mathbf{x}_r(\mathbf{s}))$, where $\mathbf{x}_r(\mathbf{s})$ is the aggregated demand in the k -DCG instance under \mathbf{s} . Therefore, with the PSNE-equivalent τ -DCG being a τ -CCG with binary demands and $\tilde{n} = \max_{j \in [\tau]} \sum_{i \in N} \tilde{d}_{i,j}$, Corollary 8 gives us the result. \square

Comparing Theorems 9 and 5, when we have the type information, Theorem 9 offers a major saving in running time by replacing $(w_{\max})^{km}$ with $\tilde{n}^{\tau m}$ in the multiplicative factor as well as $(w_{\max})^{km}$ with \tilde{n}^m (note the exponential saving of k) in the interior expression. These savings are especially pronounced when \tilde{n} is small.




Figure 2: Running-time comparison among table-based DP (TDP), set-based DP (SDP), and table-based DP asymptotic (TDPA). Encouragingly, even at small scales, TDPA hugely overestimates the actual running time. ($m = 4$ and $k = 2$.)

Theorem 9 can be extended to general k -DCGs *without* any player types, in which case $\tilde{n} = n$. This insight helps us avoid potentially large $w_{\max} \gg n$ in the running time of Theorem 5 by using Theorem 9 instead. Further running time reduction for the case of $\tilde{n} = n$ is possible through the alternative analysis given in the proof of Theorem 6.

Experiments

We have performed experiments to investigate the practical aspects of the CSP framework for non-monotonic k -DCGs with binary demand vectors.¹ Even with small-scale experiments, we show that the theoretical running-time greatly overestimates the practical, worst-case running time.

We have implemented two instantiations of the framework: (1) Table-based DP (TDP), where we use bit vectors to implement the tables, and (2) Set-based DP (SDP), where we use hash-set data structures to represent the tables. In addition, we have implemented the brute-force (BF) algorithm mentioned for the CSP shown in Fig. 1(a). BF is the only prior algorithm known to us for general k -DCGs.

All three algorithms exhaustively search for all PSNE and discard a strategy profile as soon as it is clear it cannot lead to a PSNE. We have benchmarked the theoretical running time in the worst case for a small table and extrapolated it for Theorem 5. We call it table-based DP asymptotic (TDPA). We have used non-monotonic k -DCGs with m parallel links and varied n for 15 repetitions (see the Appendix).

Fig. 2 shows that the asymptotic running time greatly overestimates the actual running time. E.g., for $n = 4$, TDPA is about eight orders of magnitude slower than SDP. Furthermore, Fig. 3 shows that SDP and TDP outperform BF easily, even for very small n . SDP is two orders of magnitude faster than BF for $n = 18$. These signify the practical appeal of our CSP framework against the backdrop of hardness re-

¹Code, unit tests, and data are in the supplementary material.




Figure 3: Running-time comparison among brute force (BF), table-based DP (TDP), and set-based DP (SDP). Even at small scales, brute force does not finish within the allocated time when $n > 20$. SDP is the fastest. ($m = 2$ and $k = 3$.)

sults. Most importantly, Procedure 2 opens up a range of possibilities for new CSP-based search algorithms rooted in, for example, backjumping and learning [Kumar, 1992, Dechter, 2003, Van Beek, 2006, Rossi et al., 2008]. We leave a comprehensive experimental study as future work.

5 LEARNING DYNAMICS APPROACH

The second class of algorithms we present is grounded in learning dynamics, which often presents a natural way of studying how players arrive at an equilibrium point [Fudenberg and Levine, 1998]. As such, learning dynamics is prominently featured in a wide range of areas from evolutionary game theory [Weibull, 1997], to wireless network [Lasaulce and Tembine, 2011], to our topic of congestion games [Shah and Shin, 2010]. In general, learning algorithms may not converge, which leads us to two threads.

First, we consider linear and exponential cost functions with convergence guarantees [Klimm and Schütz, 2022]. To our knowledge, *we are the first to derive explicit running times for k -DCGs and their variants for these cost functions*. Second, we consider the general (potentially non-monotonic) cost functions with no convergence guarantees. We present approximation algorithms for this general case.

Linear Cost Functions

We study an iterative best-response algorithm, where players start with an arbitrary strategy profile and iteratively play best responses until convergence to a PSNE, for k -DCGs and their variants using potential functions. Let the linear cost function of any resource r under a strategy profile \mathbf{s} be $c_r(\mathbf{x}_r(\mathbf{s})) \equiv a_r \sum_{j \in [k]} z_j \mathbf{x}_{r,j}(\mathbf{s}) + b_r = a_r[\mathbf{z} \cdot \mathbf{x}_r(\mathbf{s})] + b_r$, where $a_r, b_r \geq 0 \forall r$ and the k -dimensional vector $\mathbf{z} \geq 0$.

We have the following results on k -DCGs and their variants. The proofs and experimental results are in the Appendix.

Theorem 10. *For linear-cost k -DCGs, the best-response*

algorithm runs in polynomial time if $\max_r a_r, \max_r b_r$, and $\frac{\max_i [\mathbf{z} \cdot \mathbf{d}_i]^2}{\min_i [\mathbf{z} \cdot \mathbf{d}_i]}$ are polynomial in n .

Theorem 11. For linear-cost k -DCGs with binary demand vectors, the best-response algorithm runs in polynomial time if the following cost function parameters are polynomial in n : $\max_r a_r, \max_r b_r$, and $\max_j z_j$.

Theorem 12. For linear-cost k -CCGs, the best-response algorithm runs in polynomial time if $\max_r a_r, \max_r b_r, \frac{\max_j z_j^2}{\min_j z_j}$, and $\frac{\max_i d_{i,l(i)}^2}{\min_i d_{i,l(i)}}$ are polynomial in n , where $l(i) \in [k]$ denotes the index of the non-zero element in \mathbf{d}_i .

Exponential Cost Functions

Below is our result for exponential cost $c_r(\mathbf{x}_r(\mathbf{s})) \equiv a_r \exp(\mathbf{z} \cdot \mathbf{x}_r(\mathbf{s})) + b_r$. Details are in the Appendix.

Theorem 13. The best-response algorithm runs in polynomial time for exponential-cost k -DCGs if $\max_r a_r$ and $\max_r b_r$ are polynomial in n and $[\mathbf{z} \cdot \mathbf{d}_N]$ is $\mathcal{O}(\log n)$.

Approximate PSNE for General Cost Functions

We now present an (α, β) -PSNE algorithm for general cost. Let $\Delta_r \equiv \max\{\max_{i \in N, \mathbf{s} \in S; r \in s_i} c_r(\mathbf{x}_r(\mathbf{s}) - \mathbf{d}_i) - c_r(\mathbf{x}_r(\mathbf{s})), 0\}$ be the maximum non-negative marginal decrease of any player for resource r . When the congestion function is nondecreasing, $\Delta_r = 0$. Otherwise, $\Delta_r > 0$. Let $\Delta_{\max} = \max_{r \in R} \Delta_r$. We obtain the following result generalizing the result in [Christodoulou et al., 2023].

Theorem 14. Every k -DCG has an (α, β) -PSNE for $\alpha = n$ and $\beta = (n - 1)m\Delta_{\max}$. Furthermore, it can be computed using an iterative algorithm that is guaranteed to converge.

In the iterative algorithm of Theorem 14, at each round, if $\pi_i(s_i, \mathbf{s}_{-i}) > n\pi_i(s'_i, \mathbf{s}_{-i}) + (n - 1)m\Delta_{\max}$ for any player i currently playing s_i , i deviates to s'_i . As the set of strategy profiles is finite, we eventually reach an (α, β) -PSNE. The result is especially useful for small Δ_{\max} (e.g., noise).

6 STRUCTURED COSTS AND DEMANDS

Our study of structured costs and demands is motivated by a variety of realistic examples of traffic congestion games, where resources represent roads. As an example of structured/ordered demands, vehicles can be ordered by their demand vectors representing width, length, weight, etc. A common example of nondecreasing cost function is more vehicles on a road means higher costs for everyone. Singleton strategies are seen in grid-patterned road networks with parallel roads to go from source to destination [Milchtaich, 2006]. We also consider structured cost functions—e.g., different types of roads have different speed limits: highways, county routes, local roads, etc.

Ordered Demand, Nondecreasing Cost, Singleton Strat. Suppose that the players can be ordered according to their

demand vectors: $\mathbf{d}_1 \geq \mathbf{d}_2 \geq \dots \geq \mathbf{d}_n$ (w.l.o.g.). Let each player i 's set of singleton strategies $S_i = \{\{r\} \mid r \in R\}$. In addition, assume that the cost functions are nondecreasing. We can compute a PSNE using the greedy best response algorithm, which orders the players from high to low demand and lets them play their best response in that order [Milchtaich, 2006]. Details are in the Appendix.

Theorem 15. For a k -DCG with ordered demand vectors, nondecreasing cost functions, and singleton-resource strategies, a PSNE can be computed in $\mathcal{O}(n \log n + nmk)$ time.

Ordered Demand, Nondecreasing Cost, Shared Strat.

We relax the assumption of singleton-resource strategies. We show that as long as the players have the same set of strategies, we can compute a PSNE efficiently using the greedy best response algorithm.

Theorem 16. For a k -DCG with ordered demand vectors, nondecreasing cost functions, and a shared set of strategies of size p , a PSNE can be computed in $\mathcal{O}(n \log n + npmk)$.

Structured Cost Functions and Singleton Strategies

In this scenario, we do not assume any ordering among the demands of the players. Instead, we assume that the cost functions are nondecreasing and that the resources are ordered by their cost functions. That is, w.l.o.g., $c_1(\mathbf{x}) \geq c_2(\mathbf{x}) \geq \dots \geq c_m(\mathbf{x})$ for any \mathbf{x} . We also assume that there are constants $\alpha_j \geq 1$ such that $c_{j-1}(\mathbf{x}) = \alpha_j c_j(\mathbf{x})$ for any resource $j > 1$ and \mathbf{x} . These assumptions mean that some resources are more costly than others and that the costs of the resources are “nicely separated.” Finally, we assume singleton-resource strategies. We get the following result.

Theorem 17. For a k -DCG with nondecreasing and structured cost functions, where there are constants $\alpha_j \geq 1$ such that $c_{j-1}(\mathbf{x}) = \alpha_j c_j(\mathbf{x})$ for any resource $j > 1$ and aggregate demand vector \mathbf{x} , and singleton-resource strategies, a PSNE can be computed in $\mathcal{O}(n \log n + nmk)$ time.

7 CONCLUSION

We have conducted a thorough computational study of k -DCGs and their variants using two different computational methods: CSP and learning dynamics. Our dual CSP framework, which has not been studied before within the extremely rich congestion games literature, holds promise for future research within and outside of congestion games. We are particularly interested in designing and implementing CSP-inspired search algorithms—such as backjumping (Gaschnig, graph-based, conflict directed, etc.) and learning algorithms [Dechter, 2003]—for network congestion games. To our knowledge, CSP-based large-scale experimental work is yet to be done on congestion games. Beyond the realm of congestion games, our key insight of decoupling players' strategies may have applications in many other game-theoretic problems.

References

- Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM (JACM)*, 55(6):1–22, 2008.
- Eitan Altman, Thomas Boulogne, Rachid El-Azouzi, Tania Jiménez, and Laura Wynter. A survey on networking games in telecommunications. *Computers & Operations Research*, 33(2):286–311, 2006.
- Deepak Nadig Anantha, Byrav Ramamurthy, Brian Bockelman, and David Swanson. Differentiated network services for data-intensive science using application-aware sdn. In *2017 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, pages 1–6, 2017. doi: 10.1109/ANTS.2017.8384105.
- George Christodoulou, Martin Gairing, Yiannis Giannakopoulos, Diogo Poças, and Clara Waldmann. Existence and complexity of approximate equilibria in weighted congestion games. *Mathematics of Operations Research*, 48(1):583–602, 2023.
- Stella C Dafermos. The traffic assignment problem for multiclass-user transportation networks. *Transportation Science*, 6(1):73–87, 1972.
- Rina Dechter. *Constraint Processing*. Morgan Kaufmann, San Francisco, USA, 2003.
- Juliane Dunkel and Andreas S Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Mathematics of Operations Research*, 33(4):851–868, 2008.
- Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th ACM Symposium on Theory of Computing - STOC 2004*, page 604, Chicago, IL, USA, 2004. ACM Press. ISBN 978-1-58113-852-8. doi: 10.1145/1007352.1007445. URL <http://portal.acm.org/citation.cfm?doid=1007352.1007445>.
- Zubair Md Fadlullah, Yousuke Nozaki, Akira Takeuchi, and Nei Kato. A survey of game theoretic approaches in smart grid. In *2011 International Conference on Wireless Communications and Signal Processing (WCSP)*, pages 1–4, China, 2011. IEEE.
- Dimitris Fotakis, Spyros Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *International Colloquium on Automata, Languages, and Programming*, pages 123–134, Spain, 2002. Springer.
- Dimitris Fotakis, Spyros Kontogiannis, and Paul Spirakis. Selfish unsplitable flows. *Theoretical Computer Science*, 348(2):226–239, December 2005. ISSN 0304-3975. doi: 10.1016/j.tcs.2005.09.024. URL <https://www.sciencedirect.com/science/article/pii/S0304397505005347>.
- Drew Fudenberg and David K Levine. *The theory of learning in games*, volume 2. MIT press, 1998.
- Martin Gairing, Thomas Lücking, Marios Mavronicolas, and Burkhard Monien. Computing Nash equilibria for scheduling on restricted parallel links. In *Proceedings of the 36th ACM Symposium on Theory of Computing - STOC 2004*, pages 613–622, Chicago, IL, USA, 2004. ACM Press.
- Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. Pure nash equilibria: Hard and easy games. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 215–230, 2003.
- Tobias Harks and Max Klimm. On the existence of pure nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3):419–436, August 2012. ISSN 0364-765X. doi: 10.1287/moor.1120.0543. URL <https://pubsonline.informs.org/doi/abs/10.1287/moor.1120.0543>. Publisher: INFORMS.
- Tobias Harks, Max Klimm, and Rolf H Möhring. Characterizing the existence of potential functions in weighted congestion games. *Theory of Computing Systems*, 49(1):46–70, 2011.
- Max Klimm and Andreas Schütz. Congestion games with higher demand dimensions. In Tie-Yan Liu, Qi Qi, and Yinyu Ye, editors, *Web and Internet Economics*, Lecture Notes in Computer Science, pages 453–459, Cham, 2014. Springer International Publishing. ISBN 978-3-319-13129-0. doi: 10.1007/978-3-319-13129-0_39.
- Max Klimm and Andreas Schütz. Equilibria in multiclass and multidimensional atomic congestion games. *Mathematics of Operations Research*, 47(4):2743–2764, 2022.
- Vipin Kumar. Algorithms for constraint-satisfaction problems: A survey. *AI magazine*, 13(1):32–32, 1992.
- Samson Lasaulce and Hamidou Tembine. *Game theory and learning for wireless networks: fundamentals and applications*. Academic Press, 2011.
- L. Libman and A. Orda. Atomic resource sharing in non-cooperative networks. In *Proceedings of INFOCOM '97*, volume 3, pages 1006–1013 vol.3, Japan, April 1997. IEEE. doi: 10.1109/INFCOM.1997.631115. ISSN: 0743-166X.

- Marios Mavronicolas, Igal Milchtaich, Burkhard Monien, and Karsten Tiemann. Congestion games with player-specific constants. In Luděk Kučera and Antonín Kučera, editors, *Mathematical Foundations of Computer Science 2007*, pages 633–644, Berlin, Heidelberg, 2007. Springer. ISBN 978-3-540-74456-6. doi: 10.1007/978-3-540-74456-6_56.
- Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, March 1996. ISSN 08998256. doi: 10.1006/game.1996.0027. URL <https://linkinghub.elsevier.com/retrieve/pii/S0899825696900275>.
- Igal Milchtaich. The equilibrium existence problem in finite network congestion games. In *Internet and Network Economics: Second International Workshop, WINE 2006, Patras, Greece, December 15-17, 2006. Proceedings 2*, pages 87–98. Springer, 2006.
- Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
- Deepak Nadig, Byrav Ramamurthy, Brian Bockelman, and David Swanson. April: An application-aware, predictive and intelligent load balancing solution for data-intensive science. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, pages 1909–1917, 2019. doi: 10.1109/INFOCOM.2019.8737537.
- Deepak Nadig, Byrav Ramamurthy, and Brian Bockelman. Snag: Sdn-managed network architecture for gridftp transfers using application-awareness. *IEEE/ACM Transactions on Networking*, 30(4):1585–1598, 2022. doi: 10.1109/TNET.2022.3150000.
- Xidong Pi, Wei Ma, and Zhen Sean Qian. A general formulation for multi-modal dynamic traffic assignment considering multi-class vehicles, public transit and parking. *Transportation Research Part C: Emerging Technologies*, 104:369–389, 2019.
- Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, December 1973. ISSN 0020-7276, 1432-1270. doi: 10.1007/BF01737559. URL <http://link.springer.com/10.1007/BF01737559>.
- Francesca Rossi, Peter Van Beek, and Toby Walsh. Constraint programming. *Foundations of Artificial Intelligence*, 3:181–211, 2008.
- Devavrat Shah and Jinwoo Shin. Dynamics in congestion games. *ACM SIGMETRICS Performance Evaluation Review*, 38(1):107–118, 2010.
- Peter Van Beek. Backtracking search algorithms. In *Foundations of artificial intelligence*, volume 2, pages 85–134. Elsevier, 2006.
- JWC Van Lint, Serge P Hoogendoorn, and Marco Schreuder. Fastlane: New multiclass first-order traffic flow model. *Transportation Research Record*, 2088(1):177–187, 2008.
- Jian Wang, Srinivas Peeta, and Xiaozheng He. Multiclass traffic assignment model for mixed traffic flow of human-driven vehicles and connected and autonomous vehicles. *Transportation Research Part B: Methodological*, 126:139–168, 2019.
- Jörgen W Weibull. *Evolutionary game theory*. MIT press, 1997.
- Koji Yamamoto. A comprehensive survey of potential game approaches to wireless networks. *IEICE Transactions on Communications*, 98(9):1804–1823, 2015.