



HAL
open science

Dynamics on Games: Simulation-Based Techniques and Applications to Routing

Thomas Brihaye, Gilles Geeraerts, Marion Hallet, Benjamin Monmege, Bruno Quoitin

► **To cite this version:**

Thomas Brihaye, Gilles Geeraerts, Marion Hallet, Benjamin Monmege, Bruno Quoitin. Dynamics on Games: Simulation-Based Techniques and Applications to Routing. 39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2019), Dec 2019, Bombay, India. 10.4230/LIPIcs.FSTTCS.2019.35 . hal-02414987

HAL Id: hal-02414987

<https://hal.science/hal-02414987>

Submitted on 16 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamics on Games: Simulation-Based Techniques and Applications to Routing

Thomas Brihaye

Université de Mons, Mons, Belgium
thomas.brihaye@umons.ac.be

Gilles Geeraerts

Université libre Bruxelles, Brussels, Belgium
gigeerae@ulb.ac.be

Marion Hallet

Université de Mons, Mons, Belgium
Université libre Bruxelles, Brussels, Belgium
marion.hallet@umons.ac.be

Benjamin Monmege 

Aix Marseille Univ, CNRS, LIS, Université de Toulon, France
benjamin.monmege@univ-amu.fr

Bruno Quoitin

Université de Mons, Mons, Belgium
bruno.quoitin@umons.ac.be

Abstract

We consider multi-player games played on graphs, in which the players aim at fulfilling their own (not necessarily antagonistic) objectives. In the spirit of evolutionary game theory, we suppose that the players have the right to repeatedly update their respective strategies (for instance, to improve the outcome w.r.t. the current strategy profile). This generates a dynamics in the game which may eventually stabilise to an equilibrium. The objective of the present paper is twofold. First, we aim at drawing a general framework to reason about the termination of such dynamics. In particular, we identify preorders on games (inspired from the classical notion of simulation between transitions systems, and from the notion of graph minor) which preserve termination of dynamics. Second, we show the applicability of the previously developed framework to interdomain routing problems.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory

Keywords and phrases games on graphs, dynamics, simulation, network

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2019.35

Related Version A full version of the paper is available at [1], <http://arxiv.org/abs/1910.00094>.

Funding This article is based upon work from COST Action GAMENET CA 16228 supported by COST (European Cooperation in Science and Technology).

Gilles Geeraerts: Supported by an ARC grant of the Fédération Wallonie-Bruxelles.

Marion Hallet: Supported by an FNRS grant and an UMONS grant of the Fonds Franeau Mobilité.

Benjamin Monmege: Partially supported by the DeLTA project (ANR-16-CE40-0007) and by ANR project Ticktac (ANR-18-CE40-0015).

Acknowledgements We thank Timothy Griffin and Marco Chiesa for fruitful discussions.

1 Introduction

Games are nowadays a well-established model to reason about several problems in computer science. In the game paradigm, several agents (called *players*) are assumed to be rational, and interact in order to reach a fixed objective. As such, games have found numerous applications,



© Thomas Brihaye, Gilles Geeraerts, Marion Hallet, Benjamin Monmege, and Bruno Quoitin; licensed under Creative Commons License CC-BY

39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2019).

Editors: Arkadev Chattopadhyay and Paul Gastin; Article No. 35; pp. 35:1–35:14



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

such as controller synthesis [14, 17] or network protocols [12]. In this paper, we are mainly concerned about *multi-player games played on graphs*, in which $n \geq 2$ players interact trying to fulfil their own objectives (which are not necessarily antagonistic to the others); and where the arena (defining the possible actions of the players) is given as a finite graph.

An example of such game is given in Figure 1, modelling an instance of an *interdomain routing problem* which is typical of the Internet. In this case, two service providers v_1 and v_2 want to route packets to a target node v_\perp through the links that are represented by the graph edges. For economical reasons, v_1 prefers to route the traffic to v_\perp through v_2 (using path c_1s_2) instead of sending them directly to v_\perp , and symmetrically for v_2 . (Assume for instance that both v_1 and v_2 are located in Europe, and that v_\perp is in America. Then, s_1 and s_2 are transatlantic links that incur a huge cost of operation for the origin nodes.) Then, assume that, initially, v_1 and v_2 route the packets through s_1 and s_2 respectively, and broadcast this information through the network. When v_1 becomes aware of the choice of v_2 , he could decide to rely on the c_1 link instead, trying to route his packets through v_2 . However, due to the asynchronous nature of the network, v_2 could decide to route through c_2 before the new choice of v_1 reaches it. Hence, the packets get blocked in a cycle $c_1c_2c_1 \dots$ and do not reach v_\perp anymore. Then, v_1 and v_2 could decide simultaneously to reverse to s_1 and s_2 respectively which brings the network to its initial state, where the same behaviour can start again. Clearly, such oscillations in the routing policies must be avoided.

This simple example illustrates the main notions we will consider in the paper. We study the notion of *dynamics in games*, which model the behaviour of the players when they repeatedly update their strategy (i.e. their choices of actions) in order to achieve a better outcome. Then, the main objectives of the paper are to *draw a general framework to reason about the termination of such dynamics* and to *show its applicability to interdomain routing problems* (as sketched above). We say that a dynamics terminates when the players converge to an *equilibrium*, i.e. a state in which they have no incentive to further update their respective strategies. Our framework is introduced in Section 3 and 4. It relies on notions of *preorders*, in particular the *simulation preorder* [11]. Simulations are usually defined on transition systems: intuitively, a system A simulates a system B if each step of B can be mimicked in A . We consider two kinds of preorders: preorders defined on *game graphs*, i.e. on the structure of the games; and simulation defined on the *dynamics*, which are useful to reason about termination (indeed, if a dynamics D_1 simulates a dynamics D_2 , and if D_1 terminates, then D_2 terminates as well). We show how the existence of a relation between *game graphs* implies the existence of a simulation between the *induced dynamics* of those games (Theorem 8, Theorem 9). This technique allows us to *check the termination of the dynamics using structural criteria about the game graph*.

The motivation of this framework comes from several examples of problems in the literature [7, 15, 9, 2] that are (sometimes implicitly) reduced to checking the termination of a dynamics in a multi-player game, and where sufficient criteria are proposed that can be expressed as the existence of a preorder between game graphs. We thus seek to unify these results, hoping that our framework will foster new applications of the game model. For instance, several sufficient conditions for termination in the network problem sketched above consist in checking that the game graph does not contain a *forbidden pattern* [7]. This containment can naturally be expressed as a preorder.

To this aim, we introduce, in Section 4 a preorder relation on game graphs, which is inspired from the classical notion of *graph minor* [10]. Intuitively, a game graph \mathcal{G}' is a minor of \mathcal{G} if \mathcal{G}' can be obtained by deleting edges and vertices from \mathcal{G} (under well-chosen conditions that are compatible with the game setting). Then, the relation “is a minor of” forms a preorder relation on game graphs and allows one to reason on the termination of dynamics (see Theorem 8 and Theorem 9).

Finally, in Section 5, we achieve our second objective, by casting questions about Interdomain Routing into our framework. Interdomain Routing is the process of constructing routes across the networks that compose the Internet. The Border Gateway Protocol (BGP), is the *de facto* standard interdomain routing protocol. As sketched in the example above, it grows a routing tree towards every destination network in a distributed manner. The example also shows that the behaviour of the BGP is naturally modelled as a game, as already pointed out before (see [5, 15] for example). In particular, checking for so-called *safety* (does the protocol always converge to a stable state?) amounts to checking termination of some dynamics. In Section 5, we formally express BGP in our game model; revisit a classical result of Sami *et al.* that we re-prove within our framework; and finally obtain a new result regarding BGP: we provide a novel necessary and sufficient condition for convergence in the restricted (yet realistic) setting where the preferences of the nodes range on the next-hop in the route only.

Due to space constraints, full proofs and some examples can be found in [1].

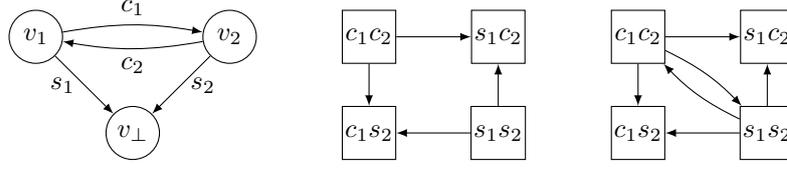
2 Preliminaries

Graphs. A (directed) *graph* is a pair $G = (V, E)$ where V is a set of *states* (or *nodes*), $E \subseteq V \times V$ is the set of *edges*. A *labelled graph* is a tuple $G = (V, E, L)$ where (V, E) is a graph, and $L : E \rightarrow S$ is a function associating, to each edge e , a label $L(e)$ from a set S of labels. A (labelled) graph G is finite iff V is finite. A *path* in a (labelled) graph G is a finite sequence $v_1 v_2 \cdots v_k$ or an infinite sequence $v_1 v_2 \cdots v_i \cdots$ of states such that $(v_i, v_{i+1}) \in E$ for all i . We denote v_1 , the first state of a path π , by $\text{first}(\pi)$. When $\pi = v_1 v_2 \cdots v_k$ is finite, we let $\text{last}(\pi) = v_k$. We let $V_\perp = \{v \in V \mid \text{there is no } v' : (v, v') \in E\}$ be the set of *terminal states*. We say that a path π is *maximal* iff: either π is infinite, or π is finite and $\text{last}(\pi) \in V_\perp$. Let $\pi_1 = v_1 \cdots v_k$ and $\pi_2 = u_1 u_2 \cdots$ be two paths such that $(v_k, u_1) \in E$. Then, we write $\pi_1 \pi_2$ to denote the new path $v_1 \cdots v_k u_1 u_2 \cdots$, obtained by the concatenation of π_1 and π_2 .

Following automata terminologies, a labelled graph G is said to be *complete deterministic* if for every state v and label ℓ , there is exactly one edge (v, v') s.t. $L(v, v') = \ell$.

Games played on graphs. An n -*player game* is a tuple $\mathcal{G} = (V, E, (V_i)_{1 \leq i \leq n}, (\preceq_i)_{1 \leq i \leq n})$ where players are denoted by $1, \dots, n$ and: (V, E) is a finite graph which forms the *arena* of the game, with V_\perp the terminal states; $(V_i)_{1 \leq i \leq n}$ is a partition of $V \setminus V_\perp$ indicating which player owns each (non-terminal) state of the game (v belongs to player i iff $v \in V_i$); and \preceq_i describes the preference of player i as a reflexive, transitive and total (i.e. for all π, π' , $\pi \preceq_i \pi'$ or $\pi' \preceq_i \pi$) binary relation defined on maximal paths which we call *plays* (the set of all plays being denoted by Play). Intuitively, player i prefers play π to play π' iff $\pi' \preceq_i \pi$. We can extract from \preceq_i a strict partial order relation by letting $\pi \prec_i \pi'$ if player i strictly prefers play π' to play π , i.e. if $\pi \preceq_i \pi'$ and $\pi' \not\preceq_i \pi$. We also write $\pi \sim_i \pi'$ if $\pi \preceq_i \pi'$ and $\pi' \preceq_i \pi$, and say that π and π' are equivalent for player i . From now on, we describe preferences by mentioning plays of interest only (implicitly, all unmentioned plays are equivalent, and below in the preference order). We also abuse notations and identify a game with its arena: so, we can write, for instance, about the “paths of \mathcal{G} ”, meaning the paths of the underlying arena.

► **Example 1.** Consider the example of [7]. In our context, it is modelled with the 2-player game $\mathcal{G}^{DIS} = (V, E, (V_1, V_2), (\preceq_1, \preceq_2))$ depicted on the left of Figure 1. The state v_\perp is terminal. Player 1 owns $V_1 = \{v_1\}$, and player 2 owns $V_2 = \{v_2\}$. Let $E = \{c_1, s_1, c_2, s_2\}$ be such that $s_i = (v_i, v_\perp)$ and $c_1 = (v_1, v_2)$, $c_2 = (v_2, v_1)$. Edges c_i stand for “continue”, and edges s_i stand for “stop”. For player 1, we let the preferences be $(v_1 v_2)^\omega \prec_1 v_1 v_\perp \prec_1 v_1 v_2 v_\perp$, where π^ω denotes an infinite number of iterations of the cycle π . Symmetrically, player 2



■ **Figure 1** Left: a 2-player game \mathcal{G}^{DIS} . Middle: $\mathcal{G}^{DIS}(\overset{P_1}{\rightarrow})$. Right: $\mathcal{G}^{DIS}(\overset{P_C}{\rightarrow})$.

has preferences $(v_2v_1)^\omega \prec_2 v_2v_\perp \prec_2 v_2v_1v_\perp$. In this case, all unmentioned plays are equally worse for both players, in particular the plays that do not start in the state owned by the player (this will always be the case in the routing application of Section 5).

Strategies and strategy profiles. The game is played by letting players move a token along the edges of the arena. Note that, in our games, there is no designated initial state, so the play can start in any state v . The choice of the initial state is not under the control of any player. Then, the player who owns v picks an edge (v, w) and moves the token to w . It is then the turn of the player who owns w to choose an edge (w, u) and so forth. The game continues *ad infinitum* or until a terminal node has been reached, thereby forming a play. Of course, each player will act in order to yield a play that is best according to his preference order \prec_i . Since no player controls the choice of the initial vertex, the players will seek to obtain the best path considering *any possible initial vertex* (see the formal definitions below). This will be important for the application of Interdomain Routing in Section 5, where the games are networks and each state corresponds to a network node that wants to send a packet to one of the terminal states.

Formally, a non-maximal path is called a *history* in the following, and the set of all histories is denoted by Hist . We let Hist_i be the set of histories h such that $\text{last}(h) \in V_i$, i.e. h ends in a state that belongs to player i . We further let $\text{player}(h) = i$ iff $h \in \text{Hist}_i$. The way players behave in the game is captured by the central notion of *strategy*, which is a mapping from a history h to a successor state in the graph, indicating how the player will play from h . A *player i strategy* is thus a function $\sigma_i: \text{Hist}_i \rightarrow V$ such that, for all $h \in \text{Hist}_i$, $(\text{last}(h), \sigma_i(h)) \in E$. A *strategy profile* σ is a tuple $(\sigma_i)_{1 \leq i \leq n}$ of strategies, one for each player i . In the following, when we consider a strategy profile σ , we always assume that σ_i is the corresponding strategy of player i . We also abuse notations, and write $\sigma(h)$ to denote the node obtained by playing the relevant strategy of σ from h , i.e. $\sigma(h) = \sigma_i(h)$ with $i = \text{player}(h)$. We denote by $\Sigma_i(\mathcal{G})$ and $\Sigma(\mathcal{G})$ the sets of player i strategies and of strategy profiles respectively (if the game \mathcal{G} is clear from the context, we may drop it and write Σ and Σ_i). As usual, given a strategy profile $\sigma = (\sigma_i)_{1 \leq i \leq n}$ and a strategy σ'_j for some player j , we denote by (σ_{-j}, σ'_j) the strategy profile obtained from σ by replacing the player j strategy σ_j with σ'_j . Fixing a history h (or, in particular, an initial node) and a profile of strategies σ is sufficient to determine a unique play that is called the *outcome*: we let $\text{Outcome}(\sigma, h)$ be the (unique) play $hv_1v_2 \dots$ such that for all $i \geq 1$: $v_i = \sigma(hv_1 \dots v_{i-1})$.

Of particular interest are the *positional strategies* (sometimes called *memoryless*), i.e. the set of strategies such that the action of the player depends on the last state of the history only. That is, σ_i is positional iff for all pairs of histories h_1 and h_2 in Hist_i : $\text{last}(h_1) = \text{last}(h_2)$ implies $\sigma_i(h_1) = \sigma_i(h_2)$. For a *positional strategy profile* σ , and a state $v \in V$, we write $\sigma(v)$ to denote the (unique) state $\sigma(h)$ returned by σ for all h with $\text{last}(h) = v$. We denote by $\Sigma^P(\mathcal{G})$ the set of strategy profiles composed of positional strategies only, and by $\Sigma_i^P(\mathcal{G})$ the set of player i positional strategies. From all states v , applying a positional strategy

profile builds a play such that the very same decision is always taken at a particular state: therefore, it either creates a finite path without cycles, or a lasso (infinite path that starts with a finite path without cycle and continues with an infinite simple cycle, disjoint from the finite path). We let Play^P be the set of all *positional plays* thus generated. In a game where we are only interested in positional strategies (as this will be the case in the application to routing, for instance), the preferences need only be defined on positional plays. Indeed, all other plays will never be obtained as an outcome, and can be assumed to be worse than any other positional play.

Game Dynamics. Let us now turn our attention to the central notion of *dynamics*. Intuitively, a dynamics consists in letting players update their strategies according to some criteria. For example, a player will want to update his strategy in order to yield a better outcome according to his preferences. Therefore, a dynamics can be understood as a graph whose states are the strategy profiles and whose edges correspond to possible updates.

► **Definition 2.** Let \mathcal{G} be a game. A dynamics for \mathcal{G} is a binary relation $\rightarrow \subseteq \Sigma \times \Sigma$ over the strategy profiles of \mathcal{G} . Its associated graph is $\mathcal{G}\langle\rightarrow\rangle = (\Sigma, \rightarrow)$, where Σ is the set of states. The terminal profiles σ of $\mathcal{G}\langle\rightarrow\rangle$ (without outgoing edges) are called the equilibria of \rightarrow .

We will focus on five dynamics, modelling certain rational behaviours of the players:

- The *one-step* dynamics $\xrightarrow{1}$. It corresponds to the minimal update that can occur, where only one player changes a single decision in order to improve the outcome from his point of view: $\sigma \xrightarrow{1} \sigma'$ iff there is a player $i \in \{1, \dots, n\}$ and a history $h \in \text{Hist}_i$ such that (i) $\sigma(h) \neq \sigma'(h)$; (ii) $\text{Outcome}(\sigma, h) \prec_i \text{Outcome}(\sigma', h)$; and (iii) $\sigma(h') = \sigma'(h')$ for all $h' \neq h$. Note that the equilibria of the one-step dynamics are exactly the so-called *subgame perfect equilibria* (SPE) introduced in [16] (see also [13]).
- The *positional one-step* dynamics $\xrightarrow{P1}$. It ranges over positional strategy profiles only, and corresponds to a single player updating his strategy from a single state. Formally, $\sigma \xrightarrow{P1} \sigma'$ (with $\sigma, \sigma' \in \Sigma^P$) iff there are a player $i \in \{1, \dots, n\}$ and a state $v \in V_i$ s.t. (i) $\sigma(v) \neq \sigma'(v)$; (ii) $\text{Outcome}(\sigma, v) \prec_i \text{Outcome}(\sigma', v)$; and (iii) $\sigma(v') = \sigma'(v')$ for all $v' \neq v$.
- The *best reply positional one-step* dynamics $\xrightarrow{bP1}$. We let $\sigma \xrightarrow{bP1} \sigma'$ iff there exists a player $i \in \{1, \dots, n\}$ and a state $v \in V_i$ such that the three properties of the positional one-step dynamics are satisfied, and, in addition, the following *best-reply* condition is satisfied: (iv) for all $\sigma'' \neq \sigma'$ such that $\sigma \xrightarrow{P1} \sigma''$ if player i is the one that has changed its strategy between σ and σ'' , then: $\text{Outcome}(\sigma'', v) \preceq_i \text{Outcome}(\sigma', v)$.
- The *positional concurrent* dynamics \xrightarrow{PC} and its best reply version \xrightarrow{bPC} . Several players can update their strategies at the same time (in a “one step” fashion), but each individual update would yield a better play when performed independently (in some sense, each player performing an update “believes” he will improve). Formally, for $\sigma, \sigma' \in \Sigma^P$, we let $\sigma \xrightarrow{PC} \sigma'$ (respectively, $\sigma \xrightarrow{bPC} \sigma'$) iff for all $i \in P(\sigma, \sigma')$, $\sigma \xrightarrow{P1} (\sigma'_i, \sigma_{-i})$ (respectively, $\sigma \xrightarrow{bP1} (\sigma'_i, \sigma_{-i})$).

Observe that other dynamics can be defined, corresponding to other behaviours of the players. We focus on these five dynamics as they fit the applications we target in Section 5. We have already said that the equilibria of $\xrightarrow{1}$ are SPEs, and we can also see from the definitions that the equilibria of the four other dynamics coincide.

► **Example 3.** Let \mathcal{G}^{DIS} be the game from Example 1. The graphs $\mathcal{G}^{\text{DIS}}\langle\xrightarrow{P1}\rangle$ and $\mathcal{G}^{\text{DIS}}\langle\xrightarrow{PC}\rangle$ are given in the middle and the right of Figure 1, where each strategy profile is represented by the choices of the players from v_1 and v_2 . For example, c_1c_2 is the strategy profile s.t.

$\sigma_1(v_1) = v_2$ and $\sigma_2(v_2) = v_1$. Note that, in this example, $\xrightarrow{P1} = \xrightarrow{bP1}$ and $\xrightarrow{PC} = \xrightarrow{bPC}$. Moreover, we can see that $\mathcal{G}^{\text{DIS}}\langle \xrightarrow{P1} \rangle$ has no infinite paths, contrary to $\mathcal{G}^{\text{DIS}}\langle \xrightarrow{PC} \rangle$. We then say that the dynamics $\xrightarrow{P1}$ terminates on \mathcal{G}^{DIS} , while \xrightarrow{PC} does not terminate on \mathcal{G}^{DIS} .

The main problem we study is whether a given dynamics terminates on a certain game: we say that a dynamics \rightarrow terminates on the game \mathcal{G} if there is no infinite path in the graph $\mathcal{G}\langle \rightarrow \rangle$ of the dynamics. As illustrated in the introduction (Example 1), such infinite paths may be problematic in certain applications, like in the Interdomain Routing problem, where an infinite path in the dynamics means that the routing protocol does not stabilise. We are thus interested in techniques to check whether a dynamics terminates on a given game.

Sometimes, a dynamics does not terminate in general, but does when we restrict ourselves to *fair executions* where all players will eventually have the opportunity to update their strategies if they want to. Formally, given a dynamics \rightarrow , an infinite path $\sigma^1 \rightarrow \sigma^2 \rightarrow \dots$ of the graph $\mathcal{G}\langle \rightarrow \rangle$ is *not fair* if there exists a player i , and a position k such that for all $\ell \geq k$, player i can switch his strategy in σ^ℓ (i.e. there is $\sigma^\ell \rightarrow \sigma'$ where $\sigma_i^\ell \neq \sigma_i'$), but for all $\ell \geq k$, player i keeps the same strategy forever (i.e. $\sigma_i^\ell = \sigma_i^k$). We say that the dynamics \rightarrow *fairly terminates* for the game \mathcal{G} if there are no infinite fair paths in the graph $\mathcal{G}\langle \rightarrow \rangle$: this is a weakening of the notion of termination seen before (see [1] for an example of a dynamics that does not terminate but terminates fairly).

3 Simulations: preorders on the dynamics graphs

At this point of the paper, it is important to understand that a game is characterised by two graphs: the *game graph* which gives its *structure* (see for example, Figure 1, left); and the *dynamics graph*, which, given a fixed dynamics \rightarrow , defines the semantics of the game as the long-term behaviour of the players (Figure 1, middle and right). In the present section, we study *preorder relations* on the *dynamics graphs*, relying on the classical notion of *simulation* [11]. They are the key ingredients to reason about the termination of dynamics.

The domain of a binary relation $R \subseteq A \times B$ is the set of elements $a \in A$ such that there exists $b \in B$ with $(a, b) \in R$. The co-domain of R is the set of elements $b \in B$ such that there exists $a \in A$ with $(a, b) \in R$. We denote the domain of R by $\text{dom}(R)$. The transitive closure R^+ of relation R is defined as $(a, b) \in R^+$ iff there are $a_0 = a, a_1, a_2, \dots, a_n = b$ such that for all $i \in \{0, 1, \dots, n-1\}$, $(a_i, a_{i+1}) \in R$.

Partial simulations and simulations. We start with some weak version of the notion of simulation, called *partial simulation* \sqsubseteq . Intuitively, we say that a state u *partially simulates* a state u' (noted $u' \sqsubseteq u$) if for all successor states v' of u' , the following holds: *if v' is in the domain of the simulation*, then there must be some state v simulating v' such that v is a successor of u . Formally, if $G = (V, E)$ and $G' = (V', E')$ are two graphs, a binary relation \sqsubseteq contained in $V' \times V$ is a *partial simulation* of G' by G if: for all $(u', v') \in E' \cap \text{dom}(\sqsubseteq)^2$, for all $u \in V$: $u' \sqsubseteq u$ implies there is $v \in V$ such that $(u, v) \in E$ and $v' \sqsubseteq v$. Then, a *simulation* \sqsubseteq of G' by G is a partial simulation of G' by G s.t. $\text{dom}(\sqsubseteq) = V'$, i.e. all states of G' are simulated by some state of G . When a (partial) simulation \sqsubseteq of G' by G exists, we say that G (partially) simulates G' . The following example highlights the difference between partial simulations and simulations. Assume G with only one edge $u \rightarrow v$ and G' with only two edges $u' \rightarrow v'_1$ and $u' \rightarrow v'_2$. Then, the relation \sqsubseteq s.t. $u' \sqsubseteq u$ and $v'_1 \sqsubseteq v$ (but $v'_2 \not\sqsubseteq v$) is a partial simulation (its domain is $\{u', v'_1\}$ so it is not a problem that v'_2 is not simulated) but is not a simulation relation.

Simulations between dynamics graphs help in showing termination properties, as shown by the following folk result:

► **Proposition 4.** *Let \mathcal{G}_1 and \mathcal{G}_2 be two games, \rightarrow_1 and \rightarrow_2 be two dynamics on \mathcal{G}_1 and \mathcal{G}_2 respectively. If $\mathcal{G}_1\langle\rightarrow_1\rangle$ simulates $\mathcal{G}_2\langle\rightarrow_2\rangle$ and the dynamics \rightarrow_1 terminates on \mathcal{G}_1 , then the dynamics \rightarrow_2 terminates on \mathcal{G}_2 .*

Bisimulations and transitive closure. We can define other preorder relations on dynamics graphs. A bisimulation is a simulation \sqsubseteq such that the inverse relation \sqsubseteq^{-1} is also a simulation. We say that $G = (V, E)$ and $G' = (V', E')$ are bisimilar when there is a bisimulation between them. As a corollary of the previous proposition, if $\mathcal{G}_1\langle\rightarrow_1\rangle$ and $\mathcal{G}_2\langle\rightarrow_2\rangle$ are bisimilar, then \rightarrow_1 terminates on \mathcal{G}_1 if and only if \rightarrow_2 terminates on \mathcal{G}_2 .

For termination purposes, it is also perfectly fine to simulate a single step of G' in several steps of G for instance. The following proposition stems from Proposition 4 and mixes the notions of transitive closures and partial simulations.

► **Proposition 5.** *Let \mathcal{G}_1 and \mathcal{G}_2 be two games, \rightarrow_1 and \rightarrow_2 be dynamics on \mathcal{G}_1 and \mathcal{G}_2 resp.*

- *If $\mathcal{G}_1\langle\rightarrow_1^+\rangle$ simulates $\mathcal{G}_2\langle\rightarrow_2\rangle$ and the dynamics \rightarrow_1 terminates on \mathcal{G}_1 , then the dynamics \rightarrow_2 terminates on \mathcal{G}_2 .*
- *If \sqsubseteq is a partial simulation of $\mathcal{G}_2\langle\rightarrow_2^+\rangle$ by $\mathcal{G}_1\langle\rightarrow_1^+\rangle$, and the dynamics \rightarrow_1 terminates on \mathcal{G}_1 , then there are no paths in $\mathcal{G}_2\langle\rightarrow_2\rangle$ that visit a state of $\text{dom}(\sqsubseteq)$ infinitely often.*

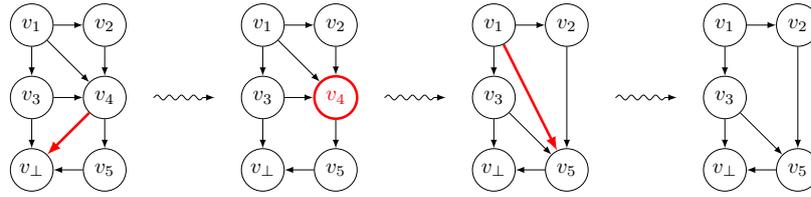
4 Minors and domination: preorders on game graphs

Let us now introduce notions of *preorders on game graphs*. We introduce a new notion of graph minor which consists in lifting the classical notion of graph minor to the context of n -player games on graphs. To the best of our knowledge, this has not been done previously. This new preorder on game graphs enables us to use in a simple context the results of Section 3 to reason about termination of dynamics. Let us start with the formal definition. For that purpose, we start by defining two transformations on game graphs. Let $\mathcal{G} = (V, E, (V_i)_i, (\preceq_i)_i)$ be an n -player game. Then we can modify it by applying either of the following transformations that yields a game $\mathcal{G}' = (V', E', (V'_i)_i, (\preceq'_i)_i)$.

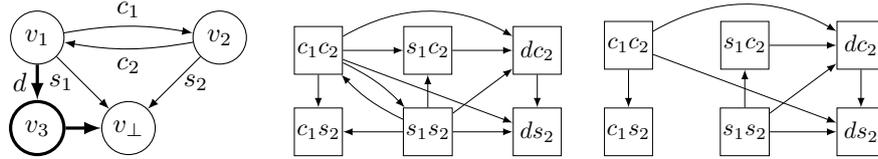
- *Deletion of an edge $(u, v) \in E$.* Then, $V' = V$, $E' = E \setminus \{(u, v)\}$, $(V'_i)_i = (V_i)_i$ and \preceq'_i is s.t. $\pi_1 \preceq'_i \pi_2$ iff $\pi_1 \preceq_i \pi_2$ and π_1, π_2 are both paths of \mathcal{G}' .
- *Deletion of a state $v \in V_j$ (for a certain player j).* This can happen in two different ways:
 1. either when v is isolated, i.e. when $(u, v) \notin E$ and $(v, u) \notin E$ for all $u \in V$. Then, $V' = V \setminus \{v\}$, $E' = E$, $V'_i = V_i$ for all $i \neq j$, $V'_j = V_j \setminus \{v\}$, and $(\preceq'_i)_i = (\preceq_i)_i$.
 2. or when v has a unique outgoing edge (v, v') and all predecessors u of v (i.e. $(u, v) \in E$) do not have v' as a successor (i.e. $(u, v') \notin E$). In this case, we have $V' = V \setminus \{v\}$, $V'_i = V_i$ for all $i \neq j$ and $V'_j = V_j \setminus \{v\}$, $E' = (E \cap (V' \times V')) \cup \{(u, v') \mid (u, v) \in E\}$, and $\pi'_1 \preceq'_i \pi'_2$ iff $\pi_1 \preceq_i \pi_2$ where π_1 and π_2 are the plays of \mathcal{G} obtained from π'_1 and π'_2 respectively, by replacing all occurrences of (u, v') (for some u) by $(u, v), (v, v')$.

► **Definition 6.** *Let \mathcal{G} and \mathcal{G}' be two n -player games. Then, \mathcal{G}' is a minor of \mathcal{G} if \mathcal{G}' can be obtained from \mathcal{G} by applying a sequence of edges and states deletions.*

► **Example 7.** An example of minor is depicted in Figure 2. If the original preferences of the player owning state v_1 are $v_1v_4v_5v_\perp \prec v_1v_3v_\perp \prec v_1v_2v_4v_\perp \prec v_1v_2v_4v_5v_\perp$ (other plays being equally worse for this player), then after the deletion of the edge (v_4, v_\perp) , his preferences become $v_1v_4v_5v_\perp \prec v_1v_3v_\perp \prec v_1v_2v_4v_5v_\perp$ (the path $v_1v_2v_4v_\perp$ does not exist in the new



■ **Figure 2** Minors obtained by first deleting the edge (v_4, v_\perp) , then the state v_4 (that has now a unique successor v_5), and then the edge (v_1, v_5) .



■ **Figure 3** L: a 3-player game \mathcal{G} with \mathcal{G}^{DIS} (Figure 1) as a minor. M: $\mathcal{G} \langle \xrightarrow{\text{PC}} \rangle$. R: $\mathcal{G} \langle \xrightarrow{\text{bPC}} \rangle$.

graph and has simply been removed from the preferences). Next, the deletion of v_4 is allowed because it is a single outgoing edge v_5 , and neither v_1 nor v_2 nor v_3 have an edge to v_5 . After this deletion, the preferences become $v_1v_5v_\perp \prec v_1v_3v_\perp \prec v_1v_2v_5v_\perp$. Finally, after the deletion of the edge (v_1, v_5) , the preferences become $v_1v_3v_\perp \prec v_1v_2v_5v_\perp$.

The deletion of a state therefore consists in squeezing each path of length 2 around it in a single edge. In the example, the deletion of the state v_4 consists in squeezing the paths $v_1v_4v_5$ in the edge v_1v_5 , and the same for $v_2v_4v_5$ and $v_3v_4v_5$ in the edges v_2v_5 and v_3v_5 respectively. The condition $(u, v') \notin E$ makes sure that this squeezing is not perturbed by the presence of an incident edge (u, v') that could have contradictory preferences. For instance, in the previous example, we cannot remove vertex v_2 in the minor obtained before having removed edge (v_1, v_5) : otherwise, we would obtain as preferences for the owner of v_1 the chain $v_1v_5v_\perp \prec v_1v_3v_\perp \prec v_1v_5v_\perp$ which is not possible.

We can link termination of dynamics on graph games to the presence of minors, in the various dynamics introduced before: if we manage to find a game minor where the dynamics does not terminate, then the original game does not terminate either.

► **Theorem 8.** *Let \mathcal{G} be a game, and \mathcal{G}' be a minor of \mathcal{G} . If $\rightarrow \in \{ \xrightarrow{1}, \xrightarrow{\text{P1}}, \xrightarrow{\text{PC}} \}$, then $\mathcal{G} \langle \rightarrow \rangle$ simulates $\mathcal{G}' \langle \rightarrow \rangle$. In particular, via Proposition 4, if the dynamics \rightarrow terminates for \mathcal{G} , then it terminates for \mathcal{G}' too.*

Sketch of proof. We prove the result for $\xrightarrow{1}$; the two other cases are similar. Since simulations are transitive relations, it is sufficient to only consider that \mathcal{G}' has been obtained from \mathcal{G} either by deleting a single edge, or by deleting a single node. Let us briefly detail the case where \mathcal{G}' is obtained by the deletion of a state v . If $h \in \text{Hist}(\mathcal{G}) \setminus \{h \mid \text{last}(h) = v\}$, we can construct a corresponding play $f(h)$ of \mathcal{G}' by replacing a sequence uvv' of h by uv' . The conditions over the deletion of v implies that that f is indeed a bijection from $\text{Hist}(\mathcal{G}) \setminus \{h \mid \text{last}(h) = v\}$ to $\text{Hist}(\mathcal{G}')$. We then consider the following relation on strategy profiles: $\sigma' \sqsubseteq \sigma$ if for all histories $h \in \text{Hist}(\mathcal{G}) \setminus \{h \mid \text{last}(h) = v\}$, $\sigma'(f(h)) = \sigma(h)$ if $\sigma(h) \neq v$, and $\sigma'(f(h)) = v'$ otherwise; and show that \sqsubseteq is a simulation (indeed a bisimulation). ◀

Notice that Theorem 8 suffers from three weaknesses. First, it does not hold for the best reply dynamics $\xrightarrow{\text{bP1}}$ and $\xrightarrow{\text{bPC}}$, as shown by the following example. Consider again the game \mathcal{G}^{DIS} from Example 1. Further, consider the game \mathcal{G} in Figure 3 obtained from \mathcal{G}^{DIS} by adding

a third player, who owns a single node v_3 , such that the only edges to and from v_3 are (v_1, v_3) and (v_3, v_\perp) , and where the preferences of player 1 are now $v_1v_\perp \prec_1 v_1v_2v_\perp \prec_1 v_1v_3v_\perp$ (observe that now, he prefers a path that traverses the new node v_3 above all other paths). Clearly, \mathcal{G}^{DIS} is a minor of \mathcal{G} . Using Theorem 8, and since we know that $\mathcal{G}^{\text{DIS}}\langle\frac{\text{PC}}{\rightarrow}\rangle$ does not terminate, we deduce that $\mathcal{G}\langle\frac{\text{PC}}{\rightarrow}\rangle$ does not terminate either. Moreover, in this example, $\mathcal{G}^{\text{DIS}}\langle\frac{\text{PC}}{\rightarrow}\rangle = \mathcal{G}^{\text{DIS}}\langle\frac{\text{bPC}}{\rightarrow}\rangle$, so even with the best-response property, the dynamics does not terminate in the minor. However, one can check that $\mathcal{G}\langle\frac{\text{bPC}}{\rightarrow}\rangle$ terminates thanks to the best-response property: Player 1 will not try to obtain path $v_1v_2v_\perp$ (which leads to a cycle in $\mathcal{G}^{\text{DIS}}\langle\frac{\text{PC}}{\rightarrow}\rangle$), but will choose a strategy going to v_3 (see Figure 3, Right). So, \mathcal{G}^{DIS} is a minor of \mathcal{G} , s.t. $\frac{\text{bPC}}{\rightarrow}$ terminates for \mathcal{G}^{DIS} but not in \mathcal{G} . The example can be adapted to $\frac{\text{bP1}}{\rightarrow}$.

A second weakness is that Theorem 8 does not apply to fair termination: the dynamics \rightarrow could fairly terminate for the game \mathcal{G} , but not for his minor \mathcal{G}' . This could be the case if we remove every choice (except one) for a certain player in the minor \mathcal{G}' creating a fair cycle in \mathcal{G}' that would not be present in \mathcal{G} .

Finally, the reciprocal of Theorem 8 does not hold: all dynamics terminate on the trivial graph with a single state, but it is also minor of all games, including those where the dynamics does not terminate.

This motivates the introduction of a stronger notion of graph minor, where it is allowed to remove *only* the so-called *dominated* edges. Formally, let \mathcal{G} be a game, let $v \in V_i$ be a state, and let $e_1 = (v, v_1)$ and $e_2 = (v, v_2)$ be two outgoing edges of v . We say that e_1 is *dominated* by e_2 if for all *positional* strategies¹ $\sigma \in \Sigma^P$, $\text{Outcome}(\sigma_1, v) \prec_i \text{Outcome}(\sigma_2, v)$, where σ_1 and σ_2 coincide with σ except that $\sigma_1(v) = v_1$ and $\sigma_2(v) = v_2$. Intuitively, this means that the player always prefers e_2 to e_1 . Then, a game \mathcal{G}' is said to be a *dominant minor* of \mathcal{G} if it can be obtained from \mathcal{G} by deleting states as before, but only deleting *dominated* edges. Equipped with this notion, we overcome the three limitations of Theorem 8 we had identified:

► **Theorem 9.** *Let \mathcal{G} be a game and \mathcal{G}' be a dominant minor of \mathcal{G} . If $\rightarrow \in \{\frac{\text{bP1}}{\rightarrow}, \frac{\text{bPC}}{\rightarrow}\}$, then we can build a simulation \sqsubseteq of $\mathcal{G}'\langle\rightarrow\rangle$ by $\mathcal{G}\langle\rightarrow\rangle$ such that: (i) \sqsubseteq^{-1} is a partial simulation of $\mathcal{G}\langle\rightarrow\rangle$ by $\mathcal{G}'\langle\rightarrow\rangle$; and (ii) if there is a fair cycle in \mathcal{G} then there is a fair cycle in \mathcal{G}' .*

In particular, the dynamics \rightarrow fairly terminates for \mathcal{G} if and only if it does for \mathcal{G}' .

Now, Theorem 9 has some limitations too. We can show that it does not hold for the “non-best-reply dynamics” $\frac{\text{P1}}{\rightarrow}$ and $\frac{\text{PC}}{\rightarrow}$. Moreover, even when we consider best-reply dynamics, the fairness condition remains crucial: we can exhibit (see [1]) a case where there is a (non-fair) cycle in \mathcal{G} but no cycles in \mathcal{G}' .

5 Applications to interdomain routing convergence

As explained in the introduction, the Border Gateway Protocol (BGP) is the *de facto* standard interdomain routing protocol. Its role is to establish routes to all the networks that compose the Internet. BGP does this by growing a routing tree towards every destination network in a distributed manner, as follows. In the initial state, only the router in the destination network has a route towards itself that it advertises to its neighbours. Each time a router receives an advertisement, it selects among the neighbour routes the one it considers best and then advertises it to its neighbours. The process repeats until no router wants to change

¹ We restrict our definition to the context of positional strategies, for the sake of brevity, but it can be extended to the more general setting.

its best route. To select its best route, a router first filters the received routes to retain only *permitted ones* and ranks them according to its *preference*. Both the filtering and ranking of routes by a router are decided based on the network's routing policy. For example, a route can be preferred over another because it offers better performance or costs less and it can be filtered out because it is not economically viable.

As shown in the introductory example, the routing approach at the heart of BGP has known convergence issues. It could fail to reach an equilibrium, entering a persistent oscillatory behaviour or it could have no equilibrium at all. This is a well-studied problem that has led to considerable work [7, 6, 5, 15, 3, 4, 8, 12]. In their seminal work [7], Griffin *et al.* analysed the BGP convergence properties using a simplified model named the Stable Path Problem (SPP). The main questions they ask are the following: (1) whether an SPP instance is **Solvable**, i.e., whether it admits a stable state; (2) whether the stable state is **Unique**; and (3) whether the system is **Safe**, i.e. it always converges to a stable state.

They also give a *sufficient* condition for an SPP instance to be safe: the absence of a substructure named a **Dispute Wheel**. Later, Sami *et al.* [15] have shown that the existence of multiple stable states is a sufficient condition to prevent safety (i.e. Safe \implies Unique). These results have later been refined by Cittadini *et al.* [3]. While the works just cited focus on the definition of sufficient conditions for safety, another approach by Gao and Rexford [6] achieves convergence by enforcing only local conditions on route preferences.

In this section, we show how SPP can be expressed in our n -player game model, therefore **Safety** reduces to checking for termination of the game dynamics. We revisit the result of Sami *et al.* by providing a new proof that relies on our framework. Then, we further exploit this framework to obtain a new result about SPP: we provide a *necessary and sufficient* condition for safety in a setting which is more restricted (yet still realistic) than in [7].

One target games. We first translate the SPP, as a combination of: (1) a reachability game that models the network topology and routing policies; and (2) the best-reply positional concurrent dynamics that models the asynchronous behaviour of the routing protocol.

Using this approach, the routing safety problem translates to a dynamics termination problem.

We rely on a particular class of games, that we call *one target games* (1TG for short): they have a unique target, the destination network, that all players want to reach. Each player corresponds to a network in the Internet and as such owns a single state. The routing policies of networks are modelled by the preference relations and by the distinction between permitted and forbidden paths. The preferences are only over positional strategies (paths), meaning that each network picks its next-hop independently of its predecessors. Permitted and forbidden paths model the fact that only some paths are allowed by the networks routing policies. Forbidden paths are also used to take into account additional restrictions that cannot be directly modelled. In SPP, the paths are simple (no loops); and non-simple paths are forbidden, for obvious reasons of efficiency. Moreover, in SPP, if at some point a network reaches a forbidden path, he will inform his neighbours that he is not able to reach the target. To model this, we impose a requirement that that if a path is permitted, all its suffixes are also permitted.

Formally, let $\mathcal{G} = (V, E, (V_i)_{1 \leq i \leq n}, (\preceq_i)_{1 \leq i \leq n})$ be an n -player game. For all $1 \leq i \leq n$, we assume that \mathcal{P}_i is the set of *permitted paths* of player i . All these paths are finite paths of the form $v_i \cdots v_\perp \in \mathcal{P}_i$. We denote by \mathcal{P}_i^c the set of *forbidden paths*, i.e. all the positional plays starting in v_i that are not in \mathcal{P}_i (in particular, all infinite paths are forbidden). We let $\mathcal{P} = \bigcup_{1 \leq i \leq n} \mathcal{P}_i$ and $\mathcal{P}^c = \bigcup_{1 \leq i \leq n} \mathcal{P}_i^c$. Then, \mathcal{G} is a *one target game* (1TG) if:

- $V_{\perp} = \{v_{\perp}\}$, and, for all players i : $V_i = \{v_i\}$;
- for all $\pi_1 \in \mathcal{P}_i^c$, for all $\pi_2 \in \mathcal{P}_i$: $\pi_1 \prec_i \pi_2$ (permitted is better than forbidden);
- for all $\pi_1, \pi_2 \in \mathcal{P}_i^c$: $\pi_1 \sim_i \pi_2$ (all forbidden paths are equivalent);
- for all $\pi_1, \pi_2 \in \mathcal{P}_i$: $\pi_1 \sim_i \pi_2$ *implies* that then there are $v \in V$ and $\tilde{\pi}_1, \tilde{\pi}_2$ s.t. $\pi_1 = v_i v \tilde{\pi}_1$ and $\pi_2 = v_i v \tilde{\pi}_2$ (if two permitted paths are equivalent, they have the same next-hop);
- for all $\pi \in \mathcal{P}_i$, for all suffixes $\tilde{\pi}$ of π : $\tilde{\pi} \in \mathcal{P}$ (all suffixes of permitted paths are permitted).

Our running example (Figure 1) is a 1TG. Since, in such a game, each player owns one and only one state, we will abuse notation by confusing each state $v \in V$ with its player. For example, for $v \in V_i$, we will write \prec_v instead of \prec_i .

Sami et al: Termination implies a unique terminal node. Equipped with this definition, we start by revisiting a result of Sami *et al.* saying that when an instance of SPP is *safe*, the solution is unique. In our setting, this translates as follows:

► **Theorem 10.** *Let \mathcal{G} be a 1TG. If $\xrightarrow{\text{bPC}}$ fairly terminates for \mathcal{G} (i.e. the corresponding instance of SPP is safe), then it has exactly one equilibrium.*

We (re-)prove this result in our setting. We rely on the notion of L -fair path that we define now. For a labelled graph $G = (V, E, L)$, we write $v_1 \rightarrow_a v_2$ iff $L(v_1, v_2) = a$ (for $v_1, v_2 \in V$). We further write $v_1 \rightarrow_A v_2$ with $A = a_1 \cdots a_n$ iff $v_1 \rightarrow_{a_1} \cdots \rightarrow_{a_n} v_2$. Then, a path $\pi = v_1 v_2 \cdots$ is L -fair if all labels occur infinitely often in this path, i.e. for all $a \in L$, for all $k \geq 1$, $\exists k' \geq k$ such that $L(v_{k'}, v_{k'+1}) = a$. A cycle π is called *constant* if there exists a state v such that $\pi = v^\omega$. Moreover, a node is a *sink* if its only outgoing edges are self loops. Then, we can show the following technical lemma:

► **Lemma 11.** *Let $G = (V, E, L)$ be a finite complete deterministic labelled graph satisfying: for all $v \in V$, for all $a, b \in L$, there are $A, B \in L^*$ and $\tilde{v} \in V$ such that $v \rightarrow_{aA} \tilde{v}$ and $v \rightarrow_{bAB} \tilde{v}$. If there exists a state from which we can reach two different sinks, then G has a non constant L -fair cycle.*

Thanks to this result, we can establish Theorem 10. We prove the contrapositive, as follows. We assume that $\mathcal{G} \langle \xrightarrow{\text{bPC}} \rangle$ has more than one equilibrium. We introduce a new dynamics \rightsquigarrow (taking into account the *beliefs* of the players about the other players' strategies) and we use Lemma 11, to show that $\mathcal{G} \langle \rightsquigarrow \rangle$ has an L -fair cycle. Then, we define a partial simulation \sqsubseteq of $\mathcal{G} \langle \rightsquigarrow \rangle$ by $\mathcal{G} \langle \xrightarrow{\text{bPC}} \rangle$ and use Proposition 5 to conclude that $\mathcal{G} \langle \xrightarrow{\text{bPC}} \rangle$ has a cycle, which is fair. Hence, $\xrightarrow{\text{bPC}}$ does not fairly terminate.

Griffin et al: Dispute wheels. Another classical notion in the BGP literature is that of a *dispute wheel* (DW for short), defined by Griffin *et al.* [7] as a “*circular set of conflicting rankings between nodes*”. They have shown that the absence of a DW is a sufficient condition for safety, which is of course of major practical interest to prove that BGP will converge in a given network. Moreover, a DW is an instance of a *forbidden pattern* in a game, and we will thus apply the results from Section 4.

We start by formally defining a DW. Let $\mathcal{G} = (V, E, (V_i)_{1 \leq i \leq n}, (\preceq_i)_{1 \leq i \leq n})$ be a 1TG with \mathcal{P}_i the set of permitted paths of v_i . A triple $D = (U, P, H)$ is a DW of \mathcal{G} if: (i) $U = (u_1, \dots, u_k) \in V^k$ is a tuple of states; (ii) $P = (\pi_1, \dots, \pi_k)$ is a tuple of permitted paths such that for all $1 \leq i \leq k$: $\pi_i \in \mathcal{P}_{u_i}$, i.e., π_i is a permitted path starting in u_i ; (iii) $H = (h_1, \dots, h_k)$ is a tuple of non-maximal paths such that for all $1 \leq i \leq k$: $h_i \pi_{(i \bmod k)+1} \in \mathcal{P}_{u_i}$; and (iv) for all $1 \leq i \leq k$: $\pi_i \prec_{u_i} h_i \pi_{(i \bmod k)+1}$.

Intuitively, in a DW, all players u_i (for $i = 1, \dots, k$) can chose between two paths to v_\perp : either a “direct” path π_i , or an “indirect” path $h_i\pi_{(i \bmod k)+1}$, which traverses $u_{(i \bmod k)+1}$; and where the latter is always preferred. So u_1 prefers to reach through u_2 , u_2 through u_3 , and so on until u_k who prefers to reach through u_1 . Such a conflict clearly yields loops where the target is never reached. The game in Figure 1 is a typical example of game that has a DW, if we let $U = (v_1, v_2)$, $P = (v_1v_\perp, v_2v_\perp)$ and $H = (v_1, v_2)$. Indeed, $v_1v_\perp \prec_1 v_1v_2v_\perp$ and $v_2v_\perp \prec_2 v_2v_1v_\perp$. Then, in our setting the sufficient condition of Griffin *et al.* [7] becomes:

► **Theorem 12** ([7]). *Let \mathcal{G} be a 1TG. If \mathcal{G} has no DW then $\xrightarrow{\text{bPC}}$ fairly terminates for \mathcal{G} .*

New result: strong dispute wheels for a necessary condition. It is well-known, however, that the absence of a DW is *not necessary* (see for example Figure 3 for a game that has a DW but where $\xrightarrow{\text{bPC}}$ terminates). As far as we know, finding a *unique and necessary* condition for the *fair* termination of $\xrightarrow{\text{bPC}}$ in 1TGs is still an open problem.

Relying on our framework, we manage to obtain such a *necessary and sufficient* condition in a restricted setting. We first strengthen the definition of DW by introducing the notion of *strong dispute wheel* (SDW for short). We then obtain two original (as far as we know) results regarding SDW. First, the absence of SDW is a *necessary* condition for the termination of $\xrightarrow{\text{PC}}$ (i.e. we drop the best-reply and the fairness hypothesis). Second, the absence of an SDW is also a *sufficient* condition in the restricted setting where the preferences of the players range only on their next-hop. This means for example that u_1 prefers to reach the target through u_2 rather than through u_3 , but does not mind the route u_2 uses (as long as v_\perp is reached). While this is a restriction, we believe that it is still meaningful in practice, since networks usually have little control about the routes chosen by their neighbours.

We first define the notion of SDW. Let \mathcal{G} be a 1TG and $D = (U, P, H)$ be a DW of \mathcal{G} . Then, D is a *strong* dispute wheel (SDW) of \mathcal{G} if:

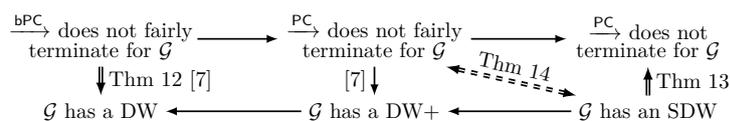
1. for all $1 \leq i \leq k$: all states $u_i \in U$ occur *only* in π_i , h_i and h_{i-1} (we identify h_0 with h_k) and not in the other paths of P and H ; and
2. for all $\pi_i, \pi_j \in P$, for all $h_k, h_\ell \in H$ with $k \neq \ell$: π_i , h_k and h_ℓ share no states of $V \setminus U$, and if π_i and π_j share a state v of $V \setminus U$ then π_i and π_j have the same suffix after v .

An important property of this definition is that, whenever a game \mathcal{G} contains an SDW $D = (U, P, H)$, we can extract a minor \mathcal{G}' which is essentially an SDW restricted to the states of U (formally, \mathcal{G}' contains an SDW $D' = (U', P', H')$ where $U' = U$ is the set of states of \mathcal{G}'). We do so by first deleting from \mathcal{G} all edges that do not occur in P and H ; then all $v \notin U$ (which have at most one outgoing edge at this point), using the procedure described in Section 4. Note that the two extra conditions in the definition of an SDW guarantee that the deletion of all the states $v \notin U$ can occur.

► **Theorem 13.** *Let \mathcal{G} be a 1TG. If $\xrightarrow{\text{PC}}$ terminates for \mathcal{G} , then \mathcal{G} has no SDW.*

Proof. By Theorem 8, it is sufficient to prove that the dynamics $\xrightarrow{\text{PC}}$ does not terminate in the minor game \mathcal{G}' extracted from the SDW (see above). We let, for all $1 \leq i \leq k$, $\sigma_1(u_i) = u_{(i \bmod k)+1}$, and $\sigma_2(u_i) = v_\perp$. Since the path resulting from σ_1 does not visit v_\perp , by definition of an SDW, we have $\sigma_1 \xrightarrow{\text{PC}} \sigma_2 \xrightarrow{\text{PC}} \sigma_1$. Hence $\mathcal{G}' \langle \xrightarrow{\text{PC}} \rangle$ contains a cycle. ◀

Thus, the absence of an SDW is a *necessary* condition for the termination of $\xrightarrow{\text{PC}}$. We can further show that this condition is *sufficient* in the restricted case where any two (permitted) paths that have the same next-hop are equivalent. Formally, let \mathcal{G} be a 1TG. We say that it



■ **Figure 4** Relationship between SDW and prior results: dashed arrow only holds for N1TG.

is a *neighbour one target game* (N1TG for short) if for all players i , for all permitted paths $\pi_1, \pi_2 \in \mathcal{P}_i$ of player i : $\pi_1 = v_i v \pi'_1$ and $\pi_2 = v_i v \pi'_2$ implies that $\pi_1 \sim_i \pi_2$. Then, we can show the following, relying on Theorem 13 (and thus, also on Theorem 8):

► **Theorem 14.** *Let \mathcal{G} be a N1TG. Then, $\frac{\text{PC}}{\rightarrow}$ does not fairly terminate for \mathcal{G} if and only if \mathcal{G} has an SDW.*

Sketch of proof. In [7], Griffin *et al.* prove a stronger result than Theorem 12, showing that if $\mathcal{G}\langle\frac{\text{PC}}{\rightarrow}\rangle$ has a fair cycle, then \mathcal{G} has a DW satisfying the following additional properties: (1) for all $u_i \in U$: $j \neq i$ implies $u_i \notin \pi_j$; (2) for all $v \notin U$, for all i, j : $v \notin \pi_i \cap h_j$; and (3) for all $v \in \pi_i \cap \pi_j$: $\pi_i(v) = \pi_j(v)$.

We call DW+ such DW. Then, the general schema of our proof is summarised in Figure 4: first, we show that the existence of a DW+ implies an SDW by showing the required additional properties. By Theorem 13, this implies $\mathcal{G}\langle\frac{\text{PC}}{\rightarrow}\rangle$ has a cycle. Then, we conclude by showing that this implies the existence of a fair cycle. ◀

Finding an SDW in practice. Because of the intricate definition of SDW, finding an SDW in a real network may be challenging in practice. However, we have:

► **Proposition 15.** *Let \mathcal{G} be an N1TG. Then $\frac{\text{PC}}{\rightarrow}$ does not fairly terminate for \mathcal{G} if and only if \mathcal{G}^{DIS} is a minor of \mathcal{G} .*

6 Perspectives

We envision multiple directions of future work. First, we could consider games with imperfect information. In the application to interdomain routing for example, this could be used to model a malicious router that advertises lies to selected neighbours. Advertising a non-existent or non-feasible path would allow for example an attacker to attract the packets of an opponent's network. Second, we could investigate a better way to model asynchronicity (useful for the routing problem) than the concurrent dynamics we have studied here. Third, we chose to model fairness via a qualitative property which ensures that *all the players will eventually have the opportunity to update their strategies if they want to*. An alternative way could be the use of probabilities: indeed, there are games for which a dynamics \rightarrow does not fairly terminate, but where an equilibrium is reached almost surely when interpreting $\mathcal{G}\langle\rightarrow\rangle$ as a finite Markov chain (with uniform distributions). Finally, we could apply the dynamics of graph-based games to other problems than interdomain routing, like *load sensitive routing*.

References

- 1 T. Brihaye, G Geeraerts, M. Hallet, Benjamin Monmege, and B. Quoitin. Dynamics on Games: Simulation-Based Techniques and Applications to Routing. *CoRR*, abs/1910.00094, 2019. [arXiv:1910.00094](https://arxiv.org/abs/1910.00094).
- 2 Thomas Brihaye, Gilles Geeraerts, Marion Hallet, and Stéphane Le Roux. Dynamics and Coalitions in Sequential Games. In Patricia Bouyer, Andrea Orlandini, and Pierluigi San Pietro, editors, *Proceedings Eighth International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2017, Roma, Italy, 20-22 September 2017.*, volume 256 of *EPTCS*, pages 136–150, 2017. [doi:10.4204/EPTCS.256.10](https://doi.org/10.4204/EPTCS.256.10).
- 3 Luca Cittadini, Giuseppe Di Battista, Massimo Rimondini, and Stefano Vissicchio. Wheel+Ring=Reel: the Impact of Route Filtering on the Stability of Policy Routing. *IEEE/ACM Transaction on Networking*, 19(4):1085–1096, August 2011.
- 4 Matthew L. Daggitt, Alexander J. T. Gurney, and Timothy G. Griffin. Asynchronous convergence of policy-rich distributed Bellman-Ford routing protocols. In Sergey Gorinsky and János Tapolcai, editors, *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM 2018, Budapest, Hungary, August 20-25, 2018*, pages 103–116. ACM, 2018. [doi:10.1145/3230543.3230561](https://doi.org/10.1145/3230543.3230561).
- 5 Alex Fabrikant and Christos H. Papadimitriou. The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In Shang-Hua Teng, editor, *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 844–853. SIAM, 2008. URL: <http://dl.acm.org/citation.cfm?id=1347082.1347175>.
- 6 Lixin Gao and Jennifer Rexford. Stable internet routing without global coordination. *IEEE/ACM Trans. Netw.*, 9(6):681–692, 2001. [doi:10.1109/90.974523](https://doi.org/10.1109/90.974523).
- 7 Timothy Griffin, F. Bruce Shepherd, and Gordon T. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Trans. Netw.*, 10(2):232–243, 2002. URL: <http://portal.acm.org/citation.cfm?id=508332>.
- 8 Aaron D. Jaggard, Neil Lutz, Michael Schapira, and Rebecca N. Wright. Dynamics at the Boundary of Game Theory and Distributed Computing. *ACM Trans. Economics and Comput.*, 5(3):15:1–15:20, 2017. [doi:10.1145/3107182](https://doi.org/10.1145/3107182).
- 9 S. Le Roux and A. Pauly. A Semi-Potential for Finite and Infinite Sequential Games (Extended Abstract). In Domenico Cantone and Giorgio Delzanno, editors, *Proceedings of the Seventh International Symposium on Games, Automata, Logics and Formal Verification*, Catania, Italy, 14-16 September 2016, volume 226 of *Electronic Proceedings in Theoretical Computer Science*, pages 242–256. Open Publishing Association, 2016. [doi:10.4204/EPTCS.226.17](https://doi.org/10.4204/EPTCS.226.17).
- 10 László Lovász. Graph minor theory. *Bull. Amer. Math. Soc. (N.S.)*, 43(1):75–86, 2006.
- 11 Robin Milner. *Communication and concurrency*. PHI Series in computer science. Prentice Hall, 1989.
- 12 Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- 13 Martin J. Osborne and Ariel Rubinstein. *A course in game theory*. MIT Press, Cambridge, MA, 1994.
- 14 A. Pnueli and R. Rosner. On the Synthesis of a Reactive Module. In *POPL*, pages 179–190. ACM Press, 1989.
- 15 Rahul Sami, Michael Schapira, and Aviv Zohar. Searching for Stability in Interdomain Routing. In *INFOCOM 2009. 28th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 19-25 April 2009, Rio de Janeiro, Brazil*, pages 549–557. IEEE, 2009. [doi:10.1109/INFCOM.2009.5061961](https://doi.org/10.1109/INFCOM.2009.5061961).
- 16 Reinhard Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrägentragheit. *Zeitschrift für die gesamte Staatswissenschaft*, 12:201–324, 1965.
- 17 Wolfgang Thomas. On the Synthesis of Strategies in Infinite Games. In *STACS*, pages 1–13, 1995.