Pseudopolynomial iterative algorithm to solve total-payoff games and min-cost reachability games

Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Benjamin Monmege

► To cite this version:

Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Benjamin Monmege. Pseudopolynomial iterative algorithm to solve total-payoff games and min-cost reachability games. *Acta Informatica*, 2017, Special Issue: Selected papers from the 26th International Conference on Concurrency Theory (CONCUR 2015), 54 (1), pp.85–125. 10.1007/s00236-016-0276-z . hal-01414114

HAL Id: hal-01414114

<https://hal.science/hal-01414114>

Submitted on 13 Dec 2016

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.

Copyright

Pseudopolynomial Iterative Algorithm to Solve Total-Payoff Games and Min-Cost Reachability Games

Thomas Brihaye · Gilles Geeraerts ·
Axel Haddad · Benjamin Monmege

30/11/2015

Abstract Quantitative games are two-player zero-sum games played on directed weighted graphs. Total-payoff games—that can be seen as a refinement of the well-studied mean-payoff games—are the variant where the payoff of a play is computed as the sum of the weights. Our aim is to describe the first pseudo-polynomial time algorithm for total-payoff games in the presence of arbitrary weights. It consists of a non-trivial application of the value iteration paradigm. Indeed, it requires to study, as a milestone, a refinement of these games, called min-cost reachability games, where we add a reachability objective to one of the players. For these games, we give an efficient value iteration algorithm to compute the values and optimal strategies (when they exist), that runs in pseudo-polynomial time. We also propose heuristics to speed up the computations.

1 Introduction

Games played on graphs are nowadays a well-studied and well-established model for the computer-aided design of computer systems, as they enable *automatic synthesis* of systems that are *correct-by-construction*. Of particular interest are *quantitative games*, that allow one to model precisely *quantitative* parameters of the system, such as energy consumption. In this setting, the game is played by two players on a directed weighted graph, where the edge weights model, for instance, a cost or a reward associated with the moves of the players. Each vertex of the graph belongs to one of the two players who compete by moving a token along the graph edges, thereby forming an infinite path called a *play*. With each play is associated a real-valued *payoff* computed from the sequence of edge weights along the

Part of this work was sponsored by EU FP7 project Cassting.

Thomas Brihaye, Axel Haddad
Université de Mons, Belgium, E-mail: thomas.brihaye,axel.haddad@umons.ac.be

Gilles Geeraerts
Université libre de Bruxelles, Belgium, E-mail: gilles.geeraerts@ulb.ac.be

Benjamin Monmege
Aix-Marseille Univ, CNRS, LIF, Marseille, France, E-mail: benjamin.monmege@lif.univ-mrs.fr

play. The traditional payoffs that have been considered in the literature include total-payoff (Gimbert and Zielonka, 2004), mean-payoff (Ehrenfeucht and Mycielski, 1979) and discounted-payoff (Zwick and Paterson, 1996). In this quantitative setting, one player aims at maximising the payoff while the other tries to minimise it. So one wants to compute, for each player, the best payoff that he can guarantee from each vertex, and the associated optimal strategies (i.e. that guarantee the optimal payoff no matter how the adversary is playing).

Such quantitative games have been extensively studied in the literature. Their associated decision problems (*is the value of a given vertex above a given threshold?*) are known to be in $\text{NP} \cap \text{co-NP}$. Mean-payoff games have arguably been best studied from the algorithmic point of view. A landmark is Zwick and Paterson’s pseudo-polynomial time (i.e. polynomial in the weighted graph when weights are encoded in unary) algorithm (Zwick and Paterson, 1996), using the *value iteration* paradigm that consists in computing a sequence of vectors of values that converges towards the optimal values of the vertices. After a fixed, pseudo-polynomial, number of steps, the computed values are precise enough to deduce the actual values of all vertices. Better pseudo-polynomial time algorithms have later been proposed, e.g., by Björklund and Vorobyov (2007); Brim et al (2011); Comin and Rizzi (2015), also achieving sub-exponential expected running time by means of randomisation.

In this paper, we focus on *total-payoff games*¹. Given an infinite play π , we denote by $\pi[k]$ the prefix of π of length k , and by $\mathbf{TP}(\pi[k])$ the (finite) sum of all edge weights along this prefix. The *total-payoff* of π , $\mathbf{TP}(\pi)$, is the inferior limit of all those sums, i.e. $\mathbf{TP}(\pi) = \liminf_{k \rightarrow \infty} \mathbf{TP}(\pi[k])$. Compared to mean-payoff (and discounted-payoff) games, the literature on total-payoff games is less extensive. Gimbert and Zielonka (2004) have shown that optimal memoryless strategies always exist for both players and the best algorithm to compute the values runs in exponential time (Gawlitza and Seidl, 2009), and consists in iteratively improving strategies. Other related works include *energy games* where one player tries to optimise its energy consumption (computed again as a sum), keeping the energy level always above 0. Note that it differs in essence from total-payoff games where no condition on the energy level is required: in particular, the optimal total-payoff could be negative, and even $-\infty$, and it is a priori not possible to simply lift all the weights by a constant to solve total-payoff games by solving a related energy games. Moreover, this difference makes difficult to apply techniques solving energy games in the case of total-payoff games. Probabilistic variants of total-payoff games have also been studied, but the weights are restricted to be non-negative (Chen et al, 2013).

We argue that the total-payoff objective is interesting as a *refinement* of the mean-payoff. Indeed, recall first that the total-payoff is finite if and only if the mean-payoff is null. Then, the computation of the total-payoff enables a finer, two-stage analysis of a game $\mathcal{G}$: (i) compute the mean payoff $\mathbf{MP}(\mathcal{G})$; (ii) subtract $\mathbf{MP}(\mathcal{G})$ from all edge weights, and scale the resulting weights if necessary to obtain integers. At that point, one has obtained a new game $\mathcal{G}'$ with null mean-payoff; (iii) compute $\mathbf{TP}(\mathcal{G}')$ to *quantify the amount of fluctuation around the mean-payoff* of the original game. Unfortunately, so far, no efficient (i.e. pseudo-polynomial time) algorithms for total-payoff games have been proposed, and straightforward

¹ Note that those games are different from *total-reward games* as studied in (Thuijsman and Vrieze, 1987).

adaptations of Zwick and Paterson’s value iteration algorithm for mean-payoff do not work, as we demonstrate at the end of Section 2. In the present article, we fill in this gap by introducing the first pseudo-polynomial time algorithm for computing the values in total-payoff games.

Our solution is a non-trivial value iteration algorithm that proceeds through nested fixed points (see Algorithm 2). A play of a total-payoff game is infinite by essence. We transform the game so that one of the players (the minimiser) must ensure a *reachability objective*: we assume that the game ends once this reachability objective has been met. The intuition behind this transformation, that stems from the use of an inferior limit in the definition of the total-payoff, is as follows: in each play π whose total-payoff is *finite*, there is a position ℓ in the play after which all the partial sums $\mathbf{TP}(\pi[i])$ (with $i \geq \ell$) will be larger than or equal to the total-payoff $\mathbf{TP}(\pi)$ of π , and infinitely often both will be equal. For example, consider the game depicted in Fig. 1(a), where the maximiser player (henceforth called **Max**) plays with the round vertices and the minimiser (**Min**) with the square vertices. For both players, the optimal value when playing from v_1 is 2, and the play $\pi = v_1 v_2 v_3 v_4 v_5 v_4 v_3 (v_4 v_5)^\omega$ reaches this value (i.e. $\mathbf{TP}(\pi) = 2$). Moreover, for all $k \geq 7$: $\mathbf{TP}(\pi[k]) \geq \mathbf{TP}(\pi)$, and infinitely many prefixes ($\pi[8]$, $\pi[10]$, $\pi[12]$, ...) have a total-payoff of 2, as shown in Fig. 1(b).

Based on this observation, we transform a total-payoff game $\mathcal{G}$, into a new game that has *the same value as the original total-payoff game* but incorporates a reachability objective for Min. Intuitively, in this new game, we allow a new action for Min: after each play prefix $\pi[k]$, he can ask to *stop the game*, in which case the payoff of the play is the payoff $\mathbf{TP}(\pi[k])$ of the prefix. However, allowing Min to stop the game at every moment would not allow us to obtain the same value as in the original total-payoff game: for instance, in the example of Fig. 1(a), Min could secure value 1 by asking to stop after $\pi[2]$, which is strictly smaller than the actual total-payoff (2) of the whole play π . So, we allow Max to *veto* to stop the game, in which case both must go on playing. Again, allowing Max to turn down all of Min’s requests would be unfair, so we parametrise the game with a natural number K , which is the maximal number of vetoes that Max can play (and we denote by $\mathcal{G}^K$ the resulting game). For the *play* depicted in Fig. 1(b), letting $K = 3$ is sufficient: trying to obtain a better payoff than the optimal, Min could request to stop after $\pi[0]$, $\pi[2]$ and $\pi[6]$, and Max can veto these three requests. After that, Max can safely accept the next request of Min, since the total payoff of all prefixes $\pi[k]$ with $k \geq 6$ are larger than or equal to $\mathbf{TP}(\pi) = 2$. Our key technical contribution is to show that *for all total-payoff games, there exists a finite, pseudo-polynomial, value of K such that the values in $\mathcal{G}^K$ and $\mathcal{G}$ coincide* (assuming all values are finite in $\mathcal{G}$: we treat the $+\infty$ and $-\infty$ values separately). Now, assume that, when Max accepts to stop the game (possibly because he has exhausted the maximal number K of vetoes), the game moves to a *target vertex*, and stops. By doing so, we effectively reduce the computation of the values in the total-payoff game $\mathcal{G}$ to the computation of the values in the total-payoff game $\mathcal{G}^K$ *with an additional reachability objective* (the target vertex) for Min.

In the following, such refined total-payoff games—where Min *must* reach a designated target vertex—will be called *min-cost reachability games*. Failing to reach the target vertices is the worst situation for Min, so the payoff of all plays that do not reach the target is $+\infty$, irrespective of the weights along the play. Otherwise, the payoff of a play is the sum of the weights up to the first occurrence of the

Fig. 1 (a) A total-payoff game, and (b) the evolution of the partial sums in π .

target. As such, this problem nicely generalises the classical shortest path problem in a weighted graph. In the one-player setting (considering the point of view of Min for instance), this problem can be solved in polynomial time by Dijkstra’s and Floyd-Warshall’s algorithms when the weights are non-negative and arbitrary, respectively. Khachiyan et al (2008) propose an extension of Dijkstra’s algorithm to handle the two-player, non-negative weights case. However, in our more general setting (two players, arbitrary weights), this problem has, as far as we know, not been studied as such, except that the associated decision problem is known to be in $\text{NP} \cap \text{co-NP}$ (Filiot et al, 2012). A pseudo-polynomial time algorithm to solve a very close problem, called the *longest shortest path problem* has been introduced by Björklund and Vorobyov (2007) to eventually solve mean-payoff games. However, because of this peculiar context of mean-payoff games, their definition of the length of a path differs from our definition of the payoff and their algorithm can not be easily adapted to solve our min-cost reachability problem. Thus, as a second contribution, we show that a value iteration algorithm enables us to compute in pseudo-polynomial time the values of a min-cost reachability game. We believe that min-cost reachability games bear their own potential theoretical and practical applications². Those games are discussed in Section 3. In addition to the pseudo-polynomial time algorithm to compute the values, we show how to compute optimal strategies for both players and characterise them: there is always a memoryless strategy for the maximiser player, but we exhibit an example (see Fig. 2) where the minimiser player needs (finite) memory. Those results on min-cost reachability games are exploited in Section 4 where we introduce and prove correct our efficient algorithm for total-payoff games.

Finally, we briefly present our implementation in Section 5, using as a core the numerical model-checker PRISM. This allows us to describe some heuristics able to improve the practical performances of our algorithms for total-payoff games and min-cost reachability games on certain subclasses of graphs.

² An example of practical application would be to perform controller synthesis taking into account energy consumption. On the other hand, the problem of computing the values in certain classes of priced timed games has recently been reduced to computing the values in min-cost reachability games (Brihaye et al, 2014).

2 Quantitative games with arbitrary weights

In this section, we formally introduce the game model we consider throughout the article.

We denote by $\mathbb{Z}$ the set of integers, and $\mathbb{Z}_\infty = \mathbb{Z} \cup \{-\infty, +\infty\}$. The set of vectors indexed by V with values in S is denoted by S^V . We let $\preceq$ be the pointwise order over $\mathbb{Z}_\infty^V$, where $x \preceq y$ if and only if $x(v) \leq y(v)$ for all $v \in V$.

2.1 Games played on graphs

We consider two-player turn-based games played on weighted graphs and denote the two *players* by Max and Min. A *weighted graph* is a tuple $\langle V, E, \omega \rangle$ where $V = V_{\text{Max}} \uplus V_{\text{Min}}$ is a finite set of vertices partitioned into the sets V_{Max} and V_{Min} of Max and Min respectively, $E \subseteq V \times V$ is a set of *directed edges*, $\omega: E \rightarrow \mathbb{Z}$ is the *weight function*, associating an integer weight with each edge. In our drawings, Max vertices are depicted by circles; Min vertices by rectangles. For every vertex $v \in V$, the set of successors of v with respect to E is denoted by $E(v) = \{v' \in V \mid (v, v') \in E\}$. Without loss of generality, we assume that every graph is deadlock-free, i.e. for all vertices v , $E(v) \neq \emptyset$. Finally, throughout this article, we let $W = \max_{(v, v') \in E} |\omega(v, v')|$ be the greatest edge weight (in absolute value) in the game graph. A *finite play* is a finite sequence of vertices $\pi = v_0 v_1 \cdots v_k \in V^*$ such that for all $0 \leq i < k$, $(v_i, v_{i+1}) \in E$. A *play* is an infinite sequence of vertices $\pi = v_0 v_1 \cdots$ such that every finite prefix $v_0 \cdots v_k$, denoted by $\pi[k]$, is a finite play.

The total-payoff of a finite play $\pi = v_0 v_1 \cdots v_k$ is obtained by summing up the weights along π , i.e. $\mathbf{TP}(\pi) = \sum_{i=0}^{k-1} \omega(v_i, v_{i+1})$. In the following, we sometimes rely on the mean-payoff to obtain information about total-payoff objectives. The *mean-payoff* computes the average weight of π , i.e. if $k \geq 1$, $\mathbf{MP}(\pi) = \frac{1}{k} \sum_{i=0}^{k-1} \omega(v_i, v_{i+1})$, and $\mathbf{MP}(\pi) = 0$ when $k = 0$. These definitions are lifted to infinite plays as follows. The total-payoff of a play π is given by $\mathbf{TP}(\pi) = \liminf_{k \rightarrow \infty} \mathbf{TP}(\pi[k])$.³ Similarly, the mean-payoff of a play π is given by $\mathbf{MP}(\pi) = \liminf_{k \rightarrow \infty} \mathbf{MP}(\pi[k])$. Tuples $\langle V, E, \omega, \mathbf{TP} \rangle$ and $\langle V, E, \omega, \mathbf{MP} \rangle$, where $\langle V, E, \omega \rangle$ is a weighted graph, are called *total-payoff* and *mean-payoff* games respectively.

2.2 Strategies and values

A *strategy* for Max (respectively, Min) in a game $\mathcal{G} = \langle V, E, \omega, \mathbf{P} \rangle$ (with $\mathbf{P}$ one of the previous payoffs), is a mapping $\sigma: V^* V_{\text{Max}} \rightarrow V$ ($\sigma: V^* V_{\text{Min}} \rightarrow V$) such that for all sequences $\pi = v_0 \cdots v_k$ with $v_k \in V_{\text{Max}}$ ($v_k \in V_{\text{Min}}$), it holds that $(v_k, \sigma(\pi)) \in E$. A play or finite play $\pi = v_0 v_1 \cdots$ conforms to a strategy σ of Max (respectively, Min) if for all k such that $v_k \in V_{\text{Max}}$ ($v_k \in V_{\text{Min}}$), we have that $v_{k+1} = \sigma(\pi[k])$. A strategy σ is *memoryless* if for all finite plays π, π' , we have that $\sigma(\pi v) = \sigma(\pi' v)$ for all $v \in V$. A strategy σ is said to be *finite-memory* if it can be encoded in a deterministic Moore machine, $\langle M, m_0, \text{up}, \text{dec} \rangle$, where M is a finite set representing the memory of the strategy, with an initial memory content $m_0 \in M$, $\text{up}: M \times V \rightarrow M$ is a

³ Our results can easily be extended by substituting a $\limsup$ for the $\liminf$. The $\liminf$ is more natural since we adopt the point of view of the maximiser Max, where the $\liminf$ is the *worst* partial sum seen infinitely often.

memory-update function, and $\text{dec}: M \times V \rightarrow V$ a decision function such that for every finite play π and vertex v , $\sigma(\pi v) = \text{dec}(\text{mem}(\pi v), v)$ where $\text{mem}(\pi)$ is defined by induction on the length of the finite play π as follows: $\text{mem}(v_0) = m_0$, and $\text{mem}(\pi v) = \text{up}(\text{mem}(\pi), v)$. In this case, we say that $|M|$ is the *size* of the strategy.

For all strategies σ_{Max} and σ_{Min} , for all vertices v , we let $\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})$ be the outcome of σ_{Max} and σ_{Min} , defined as the unique play conforming to σ_{Max} and σ_{Min} and starting in v . Naturally, the objective of Max is to maximise its payoff. In this model of zero-sum game, Min then wants to minimise the payoff of Max. Formally, we let $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Max}})$ and $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Min}})$ be the respective values of the strategies, defined as (recall that $\mathbf{P}$ is either $\mathbf{TP}$ or $\mathbf{MP}$): $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Max}}) = \inf_{\sigma_{\text{Min}}} \mathbf{P}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}}))$ and $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Min}}) = \sup_{\sigma_{\text{Max}}} \mathbf{P}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}}))$. Finally, for all vertices v , we let $\underline{\text{Val}}_{\mathcal{G}}(v) = \sup_{\sigma_{\text{Max}}} \text{Val}_{\mathcal{G}}(v, \sigma_{\text{Max}})$ and $\overline{\text{Val}}_{\mathcal{G}}(v) = \inf_{\sigma_{\text{Min}}} \text{Val}_{\mathcal{G}}(v, \sigma_{\text{Min}})$ be respectively the *lower* and *upper values* of v . We may easily show that $\underline{\text{Val}}_{\mathcal{G}} \preceq \overline{\text{Val}}_{\mathcal{G}}$. We say that strategies σ_{Max}^* of Max and σ_{Min}^* of Min are optimal if, for all vertices v : $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Max}}^*) = \underline{\text{Val}}_{\mathcal{G}}(v)$ and $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Min}}^*) = \overline{\text{Val}}_{\mathcal{G}}(v)$ respectively. We say that a game $\mathcal{G}$ is *determined* if for all vertices v , its lower and upper values are equal. In that case, we write $\text{Val}_{\mathcal{G}}(v) = \underline{\text{Val}}_{\mathcal{G}}(v) = \overline{\text{Val}}_{\mathcal{G}}(v)$, and refer to it as the *value* of v in $\mathcal{G}$. If the game is clear from the context, we may drop the index $\mathcal{G}$ from all previous notations. Mean-payoff and total-payoff games are known to be determined, with the existence of optimal memoryless strategies (Zwick and Paterson, 1996; Gimbert and Zielonka, 2004).

2.3 Previous works and contribution

Total-payoff games have been mainly considered as a refinement of mean-payoff games (Gimbert and Zielonka, 2004). Indeed, if the mean-payoff value of a game is positive (respectively, negative), its total-payoff value is necessarily $+\infty$ ($-\infty$). When the mean-payoff value is 0 however, the total-payoff is necessarily different from $+\infty$ and $-\infty$, hence total-payoff games are particularly useful in this case, to refine the analysis of the game. Deciding whether the total-payoff value of a vertex is positive can be achieved in $\text{NP} \cap \text{co-NP}$. Gawlitza and Seidl (2009) refined the complexity to $\text{UP} \cap \text{co-UP}$, and values are shown to be effectively computable solving nested fixed point equations with a strategy iteration algorithm working in exponential time in the worst case. Because of this strong relationship between mean-payoff games and total-payoff games, we can show that total-payoff games are, in some sense, as hard as mean-payoff games, for which the existence of a (strongly) polynomial time algorithm is a long-standing open question.

In this article, we improve on this state-of-the-art and introduce the first (to the best of our knowledge) pseudo-polynomial time algorithm for total-payoff games. In many cases, (e.g., mean-payoff games), a successful way to obtain such an efficient algorithm is the *value iteration paradigm*. Intuitively, value iteration algorithms compute successive approximations $x_0, x_1, \dots, x_i, \dots$ of the game value by restricting the number of turns that the players are allowed to play: x_i is the vector of optimal values achievable when the players play at most i turns. The sequence of values is computed by means of an operator $\mathcal{F}$, letting $x_{i+1} = \mathcal{F}(x_i)$ for all i . Good properties (Scott-continuity and monotonicity) of $\mathcal{F}$ ensure convergence towards its smallest or greatest fixed point (depending on the value of x_0), which, in some cases, is the value of the game.

Let us briefly explain why, unfortunately, a straightforward application of this approach fails with total-payoff games. In our case, the most natural operator $\mathcal{F}$ is such that $\mathcal{F}(x)(v) = \max_{v' \in E(v)} (\omega(v, v') + x(v'))$ for all $v \in V_{\text{Max}}$ and $\mathcal{F}(x)(v) = \min_{v' \in E(v)} (\omega(v, v') + x(v'))$ for all $v \in V_{\text{Min}}$. Indeed, this definition matches the intuition that x_N is the optimal value after N turns. Then, consider the example of Fig. 1(a), limited to vertices $\{v_3, v_4, v_5\}$ for simplicity. Observe that there are two simple cycles with weight 0, hence the total-payoff value of this game is finite. Max has the choice between cycling into one of these two cycles. It is easy to check that Max's optimal choice is to enforce the cycle between v_4 and v_5 , securing a payoff of -1 from v_4 (because of the $\liminf$ definition of **TP**). Hence, the values of v_3 , v_4 and v_5 are respectively 1, -1 and 0. In this game, we have $\mathcal{F}(x) = (2 + x(v_4), \max(-2 + x(v_3), -1 + x(v_5)), 1 + x(v_4))$, and the vector $(1, -1, 0)$ is indeed a fixed point of $\mathcal{F}$. However, it is neither the greatest nor the smallest fixed point of $\mathcal{F}$. Indeed, it is easy to check that, if x is a fixed point of $\mathcal{F}$, then $x + (a, a, a)$ is also a fixed point, for all constant $a \in \mathbb{Z} \cup \{-\infty, +\infty\}$. If we try to initialise the value iteration algorithm with value $(0, 0, 0)$, which could seem a reasonable choice, the sequence of computed vectors is: $(0, 0, 0)$, $(2, -1, 1)$, $(1, 0, 0)$, $(2, -1, 1)$, $(1, 0, 0)$, $\dots$ that is not stationary, and does not even contain $(1, -1, 0)$. Notice that $(-\infty, -\infty, -\infty)$ and $(+\infty, +\infty, +\infty)$ are fixed points, so that they do not allow us to find the correct answer too. Thus, it seems difficult to compute the actual game values with an iterative algorithm relying on the operator $\mathcal{F}$, as in the case of mean-payoff games.⁴ Notice that, in the previous example, the Zwick and Paterson's algorithm (Zwick and Paterson, 1996) to solve mean-payoff games would easily conclude from the sequence above, since the vectors of interest are then $(0, 0, 0)$, $(1, -0.5, 0.5)$, $(0.33, 0, 0)$, $(0.5, -0.25, 0.25)$, $(0.2, 0, 0)$, $\dots$ indeed converging towards $(0, 0, 0)$, the mean-payoff values of this game.

Instead, as explained in the introduction, we propose a different approach that consists in reducing total-payoff games to min-cost reachability games where Min must enforce a reachability objective on top of his optimisation objective. The aim of the next section is to study these games, and we reduce total-payoff games to them in Section 4.

3 Min-cost reachability games

In this section, we consider *min-cost reachability games* (MCR games for short), a variant of total-payoff games where one player has a reachability objective that he must fulfil first, before minimising his quantitative objective (hence the name *min-cost reachability*). Without loss of generality, we assign the reachability objective to player Min, as this will make our reduction from total-payoff games easier to explain. Hence, when the target is not reached along a path, the payoff of this path shall be the worst possible for Min, i.e. $+\infty$. Formally, an MCR game is played on a weighted graph $\langle V, E, \omega \rangle$ equipped with a target set of vertices $T \subseteq V$. The payoff $T\text{-MCR}(\pi)$ of a play $\pi = v_0 v_1 \dots$ is given by $T\text{-MCR}(\pi) = +\infty$ if the play avoids

⁴ In the context of stochastic models like Markov decision processes, Strauch (1966) already noticed that in the presence of arbitrary weights, the value iteration algorithm does not necessarily converge towards the accurate value: Puterman (1994) gives a more detailed explanation in Ex. 7.3.3.

Fig. 2 An MCR game (with W a positive integer and, v_3 the target) where Min needs memory to achieve its optimal strategy.

T , i.e. if for all $k \geq 0$, $v_k \notin T$, and $T\text{-MCR}(\pi) = \mathbf{TP}(\pi[k])$ if k is the least position in π such that $v_k \in T$. Lower and upper values are then defined as in Section 2.

Using an indirect consequence of Martin's theorem (Martin, 1975), we can show that MCR games are *determined*, i.e. that the upper and lower values always coincide:

Theorem 1 *MCR games are determined.*

Proof Consider a quantitative game $\mathcal{G} = \langle V, E, \omega, \mathbf{P} \rangle$ and a vertex $v \in V$. We will prove the determinacy result by using the Borel determinacy result of Martin (1975). First, notice that the payoff mapping $T\text{-MCR}$ is Borel measurable since the set of plays with finite $T\text{-MCR}$ payoff is a countable union of cylinders. Then, for an integer M , consider Win_M to be the set of plays with a payoff less than or equal to M . It is a Borel set, so that the qualitative game defined over the graph $\langle V, E, \omega \rangle$ with winning condition Win_M is determined. We now use this preliminary result to show our determinacy result.

We fix an MCR game and one of its vertices v , and first consider cases where either the lower or the upper values is infinite. Suppose first that $\underline{\text{Val}}(v) = -\infty$. We have to show that $\overline{\text{Val}}(v) = -\infty$ too. Let M be an integer. Since $\underline{\text{Val}}(v) < M$, we know that for all strategies σ_{Max} of Max, there exists a strategy σ_{Min} for Min, such that $\mathbf{P}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})) \leq M$. In particular, Max has no winning strategy in the qualitative game equipped with Win_M as a winning condition, hence, by determinacy, Min has a winning strategy, i.e. a strategy σ_{Min} such that every strategy σ_{Max} of Max verifies $\mathbf{P}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})) \leq M$. This exactly means that $\overline{\text{Val}}(v) \leq M$. Since this holds for every value M , we get that $\overline{\text{Val}}(v) = -\infty$. The proof goes exactly in a symmetrical way to show that $\overline{\text{Val}}(v) = +\infty$ implies $\underline{\text{Val}}(v) = +\infty$.

Consider then the case where both $\overline{\text{Val}}(v)$ and $\underline{\text{Val}}(v)$ are finite values. For the sake of contradiction, assume that $\underline{\text{Val}}(v) < \overline{\text{Val}}(v)$ and consider a real number r strictly in-between those two values. From $r < \overline{\text{Val}}(v)$, we deduce that Min has no winning strategy from v in the qualitative game with winning condition Win_r . Identically, from $\underline{\text{Val}}(v) < r$, we deduce that Max has no winning strategy from v in the same game. This contradicts the determinacy of this qualitative game. Hence, $\underline{\text{Val}}(v) = \overline{\text{Val}}(v)$. $\square$

Example 2 As an example, consider the MCR game played on the weighted graph of Fig. 2, where W is a positive integer and v_3 is the target. We claim that the values of vertices v_1 and v_2 are both $-W$. Indeed, consider the following strategy for Min: during each of the first W visits to v_2 (if any), go to v_1 ; else, go to v_3 . Clearly, this strategy ensures that the target will eventually be reached, and that either (i) edge (v_1, v_3) (with weight $-W$) will eventually be traversed; or (ii) edge (v_1, v_2) (with weight -1) will be traversed at least W times. Hence, in all plays

following this strategy, the payoff will be at most $-W$. This strategy allows Min to secure $-W$, but he can not ensure a lower payoff, since Max always has the opportunity to take the edge (v_1, v_3) (with weight $-W$) instead of cycling between v_1 and v_2 . Hence, Max's optimal choice is to follow the edge (v_1, v_3) as soon as v_1 is reached, securing a payoff of $-W$. The Min strategy we have just given is optimal, and there is *no optimal memoryless strategy* for Min. Indeed, always playing (v_2, v_3) does not ensure a payoff $\leq -W$; and, always playing (v_2, v_1) does not guarantee to reach the target, and this strategy has thus value $+\infty$.

A remark on related work Let us note that Björklund and Vorobyov (2007) introduce the *longest shortest path problem* (LSP for short) and propose a pseudo-polynomial time algorithm to solve it. However, their definition has several subtle but important differences to ours, such as in the definition of the payoff of a play (equivalently, the length of a path). As an example, in the game of Fig. 2, the play $\pi = (v_1 v_2)^\omega$ (that never reaches the target) has length $-\infty$ in their setting, while, in our setting, $\{v_3\}\text{-MCR}(\pi) = +\infty$. A more detailed comparison of the two definitions is given in Appendix A. Moreover, even if a preprocessing would hypothetically allow one to use the LSP algorithm to solve MCR games, our solution (that has the same worst-case complexity as theirs) is simpler to implement, and we also introduce (see Section 5) heuristics that are only applicable to our value iteration solution.

As explained in the introduction of this section, we show how to *solve* those games, i.e. how to compute $\text{Val}(v)$ for all vertices v in pseudo-polynomial time. This procedure will be instrumental to solving total-payoff games. Our contributions are summarised in the following theorem:

Theorem 3 *Let $\mathcal{G} = \langle V, E, \omega, T\text{-MCR} \rangle$ be an MCR game.*

1. *For all $v \in V$, deciding whether $\text{Val}(v) = +\infty$ can be done in polynomial time.*
2. *For $v \in V$, deciding whether $\text{Val}(v) = -\infty$ is as hard as solving mean-payoff games, in $\text{NP} \cap \text{co-NP}$ and can be achieved in pseudo-polynomial time.*
3. *If $\text{Val}(v) \neq -\infty$ for all vertices $v \in V$, then both players have optimal strategies. Moreover, Max always has a memoryless optimal strategy, while Min may require finite (pseudo-polynomial) memory in his optimal strategy.*
4. *Computing all values $\text{Val}(v)$ (for $v \in V$), as well as optimal strategies (if they exist) for both players, can be done in (pseudo-polynomial) time $O(|V|^2|E|W)$.*

3.1 Finding vertices with value $+\infty$

To prove the first item of Theorem 3, it suffices to notice that vertices with value $+\infty$ are exactly those from which Min can not reach the target. Therefore the problem reduces to deciding the winner in a classical reachability game, that can be solved in polynomial time (Thomas, 1995), using the classical *attractor* construction.

More precisely, let $\mathcal{G} = \langle V, E, \omega, T\text{-MCR} \rangle$ be an MCR game. Notice that for all plays $\pi = v_0 v_1 \dots$, $T\text{-MCR}(\pi) = +\infty$ if and only if $v_k \notin T$ for all $k \geq 0$, i.e. π avoids the target. Then, let us show that the classical attractor technique (Thomas, 1995) allows us to compute the set $V_{+\infty} = \{v \in V \mid \text{Val}(v) = +\infty\}$.

Recall that the attractor of a set T of vertices is obtained thanks to the sequence $\text{Attr}_0(T), \dots, \text{Attr}_i(T), \dots$ where: $\text{Attr}_0(T) = T$; and for all $i \geq 0$:

$$\begin{aligned} \text{Attr}_{i+1}(T) = & \text{Attr}_i(T) \cup \{v \in V_{\text{Min}} \mid E(v) \cap \text{Attr}_i(T) \neq \emptyset\} \\ & \cup \{v \in V_{\text{Max}} \mid E(v) \subseteq \text{Attr}_i(T)\}. \end{aligned}$$

It is well-known that this sequence converges after at most $|V|$ steps to the set $\text{Attr}(T)$ of all vertices from which Min has a memoryless strategy to ensure reaching T . Hence, under our hypothesis, $V_{+\infty} = V \setminus \text{Attr}(T)$. This proves the first item of Theorem 3. Observe that we can safely *remove* from the game graph all vertices v such that $\text{Val}(v) = +\infty$, without changing the values of the other vertices. Hence, we can, when need be, assume that the MCR games we consider contain no vertex with value $+\infty$, as they can be removed by this polynomial-time preprocessing.

In those games, one can construct in polynomial time a memoryless strategy, called an *attractor strategy*, ensuring to reach the target in less than $|V|$ steps from every vertex.

In the following, we assume that all vertices have a value different from $+\infty$. Indeed as described above one can detect in polynomial time the vertices with value $+\infty$ and remove them *without changing the values of the other vertices*.

3.2 Finding vertices with value $-\infty$

To prove the second item, we notice that vertices with value $-\infty$ are exactly those with a value < 0 in the mean-payoff game played on the same graph. On the other hand, we can show that every mean-payoff game can be transformed (in polynomial time) into an MCR game such that a vertex has value < 0 in the mean-payoff game if and only if the value of its corresponding vertex in the MCR game is $-\infty$. More precisely:

Proposition 4

1. For all MCR games $\mathcal{G} = \langle V, E, \omega, T\text{-MCR} \rangle$ where $\text{Val}_{\mathcal{G}}(v) \neq +\infty$ for all v , for all vertices v of $\mathcal{G}$, $\text{Val}_{\mathcal{G}}(v) = -\infty$ if and only if $\text{Val}_{\mathcal{G}'}(v) < 0$, where $\mathcal{G}'$ is the mean-payoff game $\langle V, E, \omega, \text{MP} \rangle$.
2. Conversely, given a mean-payoff game $\mathcal{G} = \langle V, E, \omega, \text{MP} \rangle$, we can build, in polynomial time, an MCR game $\mathcal{G}'$ such that for all vertices v of $\mathcal{G}$: $\text{Val}_{\mathcal{G}}(v) < 0$ if and only if $\text{Val}_{\mathcal{G}'}(v) = -\infty$.

Proof To prove the first item, consider an MCR game $\mathcal{G} = \langle V, E, \omega, T\text{-MCR} \rangle$ such that $\text{Val}_{\mathcal{G}}(v) \neq +\infty$ for all $v \in V$, and $\mathcal{G}' = \langle V, E, \omega, \text{MP} \rangle$ the same weighted graph equipped with a mean-payoff objective.

If $\text{Val}_{\mathcal{G}'}(v) < 0$, we know that there is a profile of optimal memoryless strategies $(\sigma_{\text{Max}}^*, \sigma_{\text{Min}}^*)$ such that the outcome starting in v and following this profile necessarily starts with a finite prefix and then loops in a cycle with a total weight less than 0. For every $M > 0$, we construct a strategy σ_{Min}^M that ensures in $\mathcal{G}$ a cost less than or equal to $-M$: this will prove that $\text{Val}_{\mathcal{G}}(v) = -\infty$. Since we have assumed that $\text{Val}_{\mathcal{G}}(v) \neq +\infty$ for all v , we know that Min has a strategy to reach the target from all v (for instance, take the attractor strategy described above), by a path of length at most $|V|$. Thus, there exists a bound w and a strategy allowing Min

to reach the target from every vertex of $\mathcal{G}'$ with a cost at most w . The strategy σ_{Min}^M of Min is then to follow σ_{Min}^* until the accumulated cost is less than $-M - w$, at which point it follows his strategy to reach the target. Clearly, for all M , σ_{Min}^M guarantees that Min reaches the target with a cost at most $-M$.

Reciprocally, if $\text{Val}_{\mathcal{G}}(v) = -\infty$, consider $M = |V|W$ and a strategy σ_{Min}^M of Min ensuring a cost less than $-M$, i.e. such that $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Min}}^M) < -M$. Consider the finitely-branching tree built from $\mathcal{G}$ by unfolding the game from vertex v and resolving the choices of Min with strategy σ_{Min}^M . Each branch of this tree corresponds to a possible strategy of Max. Since this strategy generates a finite cost, we are certain that every such branch leads to a vertex of T . If we trim the tree at those vertices, we finally obtain a finite tree. Now, for a contradiction, consider an optimal memoryless strategy σ_{Max}^* of Max securing a non-negative mean-payoff, that is, $\text{Val}_{\mathcal{G}'}(v, \sigma_{\text{Max}}^*) \geq 0$. Consider the branch of the previous tree where Max follows strategy σ_{Max}^* . Since this finite branch has cost less than $-M = -|V|W < 0$ (W is positive, otherwise the mean-payoff value would be 0), we know for sure that there are two occurrences of the same vertex v' with an in-between weight less than 0: otherwise, by removing all non-negative cycles, we obtain a play without repetition of vertices, henceforth of length bounded by $|V|$, and therefore of cost at least $-M$. Suppose that $v' \in V_{\text{Max}}$. Then, Min has a strategy σ_{Min} to ensure a negative mean-payoff $\text{Val}_{\mathcal{G}'}(v, \sigma_{\text{Min}}) < 0$: indeed, he simply modifies⁵ his strategy so that he always stays in the negative cycle starting in v' (he can do that since σ_{Max} plays a memoryless strategy, so that he can not change his decisions in the cycle), ensuring that, against the optimal strategy σ_{Max}^* of Max, he gets a mean-payoff being the cost of the cycle. This is a contradiction since Max is supposed to have a strategy ensuring a non-negative mean-payoff from v . Hence, $v' \in V_{\text{Min}}$. But the same contradiction appears in that case since Min can force that it always stays in the negative cycle by modifying his strategy. Finally, we have proved that Max can not have a memoryless strategy securing a non-negative mean-payoff from v . By memoryless determinacy of the mean-payoff games, this ensures that Min has a memoryless strategy securing a negative mean-payoff from v .

Hence, we have shown that $\text{Val}_{\mathcal{G}}(v) = -\infty$ if and only if $\text{Val}_{\mathcal{G}'}(v) < 0$, which concludes the first claim of Proposition 4.

To prove the second item, we reduce mean-payoff games to MCR games as follows. Let $\mathcal{G} = \langle V, E, \omega, \mathbf{MP} \rangle$ be a mean-payoff game. Without loss of generality, we may suppose that the graph of the game is bipartite, in the sense that $E \subseteq V_{\text{Max}} \times V_{\text{Min}} \cup V_{\text{Min}} \times V_{\text{Max}}$.⁶ The problem we are interested in is to decide whether $\text{Val}_{\mathcal{G}}(v) < 0$ for a given vertex v . We now construct an MCR game $\mathcal{G}' = \langle V', E', \omega', T', \mathbf{MCR} \rangle$ from $\mathcal{G}$. The only difference is the presence of a fresh target vertex $\mathbf{t}$ on top of vertices of V : $V' = V \uplus \{\mathbf{t}\}$ with $T' = \{\mathbf{t}\}$. Edges of $\mathcal{G}'$ are given by $E' = E \cup \{(v, \mathbf{t}) \mid v \in V_{\text{Min}}\} \cup \{(\mathbf{t}, \mathbf{t})\}$. Weights of edges are given by: $\omega'(v, v') = \omega(v, v')$ if $(v, v') \in E$, and $\omega'(v, \mathbf{t}) = \omega'(\mathbf{t}, \mathbf{t}) = 0$. We show that $\text{Val}_{\mathcal{G}}(v) < 0$ if and only if $\text{Val}_{\mathcal{G}'}(v) = -\infty$.

⁵ This is not needed in the proof, but notice that Min necessarily modifies its strategy here, i.e. owns at least one vertex of the cycle. Otherwise, the value in the MCR game $\mathcal{G}$ would not be $-\infty$.

⁶ It suffices to add a vertex of the opponent in-between two vertices of the same player related by a transition in $\mathcal{G}$: in this vertex, the opponent has no choice but to follow the transition chosen by the first player.

Algorithm 1: Value iteration for MCR games. Gray lines correspond to the computation of optimal strategy for both players (see Section 3.4)

Input: MCR game $\langle V, E, \omega, \mathbf{MCR} \rangle$, W greatest weight in absolute value

```

1  $X(\mathbf{t}) := 0$ 
2 foreach  $v \in V \setminus \{\mathbf{t}\}$  do  $X(v) := +\infty$ 
3 repeat
4    $X_{pre} := X$ 
5   foreach  $v \in V_{\text{Max}} \setminus \{\mathbf{t}\}$  do
6      $X(v) := \max_{v' \in E(v)} (\omega(v, v') + X_{pre}(v'))$ 
7      $\sigma_{\text{Max}}^*(v) := \operatorname{argmax}_{v' \in E(v)} (\omega(v, v') + X_{pre}(v'))$ 
8   foreach  $v \in V_{\text{Min}} \setminus \{\mathbf{t}\}$  do
9      $X(v) := \min_{v' \in E(v)} (\omega(v, v') + X_{pre}(v'))$ 
10    if  $X(v) \neq X_{pre}(v)$  then
11       $\sigma_{\text{Min}}^*(v) := \operatorname{argmin}_{v' \in E(v)} (\omega(v, v') + X_{pre}(v'))$ 
12      if  $X_{pre}(v) = +\infty$  then  $\sigma_{\text{Min}}^\dagger(v) = \sigma_{\text{Min}}^*(v)$ 
13  foreach  $v \in V \setminus \{\mathbf{t}\}$  do
14    if  $X(v) < -( |V| - 1 )W$  then  $X(v) := -\infty$ 
15 until  $X = X_{pre}$ 
16 return  $X$ 

```

In $\mathcal{G}'$, all values are different from $+\infty$, since Min plays at least every two steps, and has the capability to go to the target vertex with weight 0. Hence, letting $\mathcal{G}'' = \langle V', E', \omega', \mathbf{MP} \rangle$ the mean-payoff game on the weighted graph of $\mathcal{G}'$, by the previous direction, we have that for every vertex $v \in V'$, $\text{Val}_{\mathcal{G}'}(v) = -\infty$ if and only if $\text{Val}_{\mathcal{G}''}(v) < 0$.

To conclude, we prove that for all vertices $v \in V$, $\text{Val}_{\mathcal{G}''}(v) < 0$ if and only if $\text{Val}_{\mathcal{G}}(v) < 0$. If $\text{Val}_{\mathcal{G}}(v) < 0$, by mapping the memoryless optimal strategies of $\mathcal{G}$ into $\mathcal{G}''$, we directly obtain that $\text{Val}_{\mathcal{G}''}(v) \leq \text{Val}_{\mathcal{G}}(v) < 0$, since Max has no possibility to go by himself to the target. Reciprocally, if $\text{Val}_{\mathcal{G}''}(v) < 0$, we can project a profile of memoryless optimal strategies over vertices of $\mathcal{G}$, since the target can not be visited in this case (otherwise the optimal play would have mean-payoff 0): the play obtained from v in $\mathcal{G}$ is then the projection of the play obtained from v in $\mathcal{G}''$, with the same cost. Hence, $\text{Val}_{\mathcal{G}}(v) \leq \text{Val}_{\mathcal{G}''}(v) < 0$. $\square$

3.3 Computing all values

Now that we have discussed the case of vertices with value in $\{-\infty, +\infty\}$, let us present our core contribution on MCR games, which is a pseudo-polynomial time, value iteration algorithm to compute the values of those games. Note that this algorithm is correct even when some vertices have value in $\{-\infty, +\infty\}$, as we will argue later.

In all that follows, we assume that there is exactly one target vertex denoted by $\mathbf{t}$, and the only outgoing edge from $\mathbf{t}$ is a self loop with weight 0: this is reflected by denoting $\mathbf{MCR}$ the payoff mapping $\{\mathbf{t}\}$ - $\mathbf{MCR}$. This is without loss of generality since everything that happens after the first occurrence of a target vertex in a play does not matter for the payoff.

Our value iteration algorithm for MCR games is given in Algorithm 1. As it can be seen, this algorithm consists in computing a sequence of vectors $\mathbf{X}$. Initially (line 2), $\mathbf{X}(v) = +\infty$ for all vertices but the target $\mathbf{t}$ where $\mathbf{X}(\mathbf{t}) = 0$. Then, a new value of $\mathbf{X}$ is obtained by optimising locally the value of each node, and changing to $-\infty$ the value $\mathbf{X}(v)$ of all vertices v such that the computed value $\mathbf{X}(v)$ has gone below a given threshold $-(|V| - 1)W$ (line 14). The following proposition states the correctness of Algorithm 1.

Proposition 5 *If an MCR game $\mathcal{G} = \langle V, E, \omega, \mathbf{MCR} \rangle$ is given as input (possibly with values $+\infty$ or $-\infty$), Algorithm 1 outputs $\text{Val}_{\mathcal{G}}$, after at most $(2|V| - 1)W|V| + 2|V|$ iterations.*

To establish this proposition, we consider the sequence of values $(x_i)_{i \geq 0}$ that vector $\mathbf{X}$ takes along the execution of the algorithm. More formally, we can define this sequence thanks to the operator $\mathcal{F}$, which denotes the function $\mathbb{Z}_{\infty}^V \rightarrow \mathbb{Z}_{\infty}^V$ mapping every vector $x \in \mathbb{Z}_{\infty}^V$ to $\mathcal{F}(x)$ defined, for all vertices v , by:

$$\mathcal{F}(x)(v) = \begin{cases} 0 & \text{if } v = \mathbf{t} \\ \max_{v' \in E(v)} (\omega(v, v') + x(v')) & \text{if } v \in V_{\text{Max}} \setminus \{\mathbf{t}\} \\ \min_{v' \in E(v)} (\omega(v, v') + x(v')) & \text{if } v \in V_{\text{Min}} \setminus \{\mathbf{t}\}. \end{cases}$$

Then, for all vertices v we let $x_0(\mathbf{t}) = 0$, and $x_0(v) = +\infty$ for all $v \neq \mathbf{t}$. Moreover, for all $i \geq 1$, we let $x_i = \mathcal{F}(x_{i-1})$.

The intuition behind this sequence is that x_i is the value of the game if we impose that Min must reach the target within i steps (and get a payoff of $+\infty$ if he fails to do so). Notice that operator $\mathcal{F}$ is monotonic (i.e. $\mathcal{F}(x) \preceq \mathcal{F}(y)$ for all $x \preceq y$), and that $x_1 = \mathcal{F}(x_0) \preceq x_0$, so that we know that the sequence $(x_i)_{i \geq 0}$ is non-increasing:

$$\forall i \geq 0 \quad x_{i+1} \preceq x_i. \quad (1)$$

In order to formalise the intuition that x_i is the value of the game if we impose that Min must reach the target within i steps, we define, for a play $\pi = v_0 v_1 \dots v_i \dots$:

$$\mathbf{MCR}^{\leq i}(\pi) = \begin{cases} \mathbf{MCR}(\pi) & \text{if } v_k = \mathbf{t} \text{ for some } k \leq i \\ +\infty & \text{otherwise.} \end{cases}$$

We further let $\overline{\text{Val}}^{\leq i}(v) = \inf_{\sigma_{\text{Min}}} \sup_{\sigma_{\text{Max}}} \mathbf{MCR}^{\leq i}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}}))$ (where σ_{Max} and σ_{Min} are respectively strategies of Max and Min). Observe first that for all $v \in V$, for all $i \geq 0$, and for all strategies σ_{Max} and σ_{Min} :

$$\mathbf{MCR}^{\leq i}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})) \geq \mathbf{MCR}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})).$$

Indeed, if the target vertex $\mathbf{t}$ is reached within i steps, then payoffs are equal. Otherwise, $\mathbf{MCR}^{\leq i}(\text{Play}(v, \sigma_{\text{Min}}, \sigma_{\text{Max}})) = +\infty$. Thus, for all $i \geq 1$ and $v \in V$:

$$\overline{\text{Val}}^{\leq i}(v) \geq \overline{\text{Val}}(v) = \text{Val}(v)$$

which can be rewritten as

$$\overline{\text{Val}}^{\leq i} \succeq \overline{\text{Val}} = \text{Val}.$$

Let us now consider the sequence $(\overline{\text{Val}}^{\leq i})_{i \geq 0}$. We first give an alternative definition of this sequence permitting to show its convergence.

Lemma 6 For all $i \geq 1$, for all $v \in V$:

$$\overline{\text{Val}}^{\leq i}(v) = \begin{cases} 0 & \text{if } v = \mathbf{t} \\ \max_{v' \in E(v)} \left(\omega(v, v') + \overline{\text{Val}}^{\leq i-1}(v') \right) & \text{if } v \in V_{\text{Max}} \setminus \{\mathbf{t}\} \\ \min_{v' \in E(v)} \left(\omega(v, v') + \overline{\text{Val}}^{\leq i-1}(v') \right) & \text{if } v \in V_{\text{Min}} \setminus \{\mathbf{t}\}. \end{cases}$$

Proof The lemma can be established by showing that $\overline{\text{Val}}^{\leq i}(v)$ is the value in a game played on a finite tree of depth i (i.e. by applying a backward induction). We adopt the following notation for labeled unordered trees. A leaf is denoted by (v) , where $v \in V$ is the label of the leaf. A tree with root labeled by v and subtrees $A_1, \dots, A_n$ is denoted by $(v, \{A_1, \dots, A_n\})$. Then, for each $v \in V$ and $i \geq 0$, we define $A^i(v)$ as follows:

$$A^0(v) = (v) \\ \text{for all } i \geq 1 : A^i(v) = (v, \{A^{i-1}(v') \mid (v, v') \in E\}).$$

Now, let us further label those trees by a value in $\mathbb{Z} \cup \{+\infty\}$ thanks to the function λ (thereby formalising the backward induction). For all trees of the form $A^0(v) = (v)$, we let:

$$\lambda(A^0(v)) = \begin{cases} 0 & \text{if } v = \mathbf{t} \\ +\infty & \text{if } v \neq \mathbf{t}. \end{cases}$$

For all trees of the form $A^i(v) = (v, \{A^{i-1}(v_1), \dots, A^{i-1}(v_m)\})$ (for some $i \geq 1$), we let:

$$\lambda(A^i(v)) = \begin{cases} 0 & \text{if } v = \mathbf{t} \\ \max_{1 \leq j \leq m} \left(\omega(v, v_j) + \lambda(A^{i-1}(v_j)) \right) & \text{if } v \in V_{\text{Max}} \setminus \{\mathbf{t}\} \\ \min_{1 \leq j \leq m} \left(\omega(v, v_j) + \lambda(A^{i-1}(v_j)) \right) & \text{if } v \in V_{\text{Min}} \setminus \{\mathbf{t}\}. \end{cases} \quad (2)$$

Clearly, for all $v \in V$, for all $i \geq 0$, the branches of $A^i(v)$ correspond to all the possible finite plays $\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})[i]$, i.e. there is a branch for each possible strategy profile $(\sigma_{\text{Max}}, \sigma_{\text{Min}})$. Thus, $\lambda(A^i(v)) = \overline{\text{Val}}^{\leq i}(v)$ for all $i \geq 0$, which permits us to conclude from (2). $\square$

We have just shown that for all $i \geq 1$, $\overline{\text{Val}}^{\leq i} = \mathcal{F}(\overline{\text{Val}}^{\leq i-1})$, and since $x_0 = \overline{\text{Val}}^{\leq 0}$, we obtain (as expected) $x_i = \overline{\text{Val}}^{\leq i}$ for all $i \geq 0$. The main question is now to characterise the limit of the sequence $(x_i)_{i \geq 0}$, and more precisely, to prove that it is the value Val of the game. Indeed, at this point, it would not be too difficult to show that Val is a fixed point of operator $\mathcal{F}$, but it would be more difficult to show that it is the *greatest* fixed point of $\mathcal{F}$, that is indeed the limit of sequence $(x_i)_{i \geq 0}$ (by Kleene's theorem, applicable since $\mathcal{F}$ is Scott-continuous). Instead, we study refined properties of the sequence $(\overline{\text{Val}}^{\leq i})_{i \geq 0}$, namely its stationarity and the speed of its convergence, and deduce that Val is the greatest fixed point as a corollary (see Corollary 11).

We start by characterising how $\overline{\text{Val}}^{\leq i}$ evolves over the first $|V| + 1$ steps. The next lemma states that, for each node v , the sequence $\overline{\text{Val}}^{\leq 0}(v), \overline{\text{Val}}^{\leq 1}(v), \dots, \overline{\text{Val}}^{\leq i}(v), \dots, \overline{\text{Val}}^{\leq |V|}(v)$ is of the form

$$\underbrace{+\infty, +\infty, \dots, +\infty}_{k \text{ times}}, a_k, a_{k+1}, \dots, a_{|V|}$$

where k is the step at which v has been added to the attractor, and each value a_i is finite and bounded:

Lemma 7 *Let $v \in V$ be a vertex and let $0 \leq k \leq |V|$ be such that $v \in \text{Attr}_k(\{\mathbf{t}\}) \setminus \text{Attr}_{k-1}(\{\mathbf{t}\})$ (assuming $\text{Attr}_{-1}(\{\mathbf{t}\}) = \emptyset$). Then, for all $0 \leq j \leq |V|$: (i) $j < k$ implies $\overline{\text{Val}}^{\leq j}(v) = +\infty$ and (ii) $j \geq k$ implies $\overline{\text{Val}}^{\leq j}(v) \leq jW$.*

Proof We prove the property for all vertices v , by induction on j .

Base case: $j = 0$. We consider two cases. Either $v = \mathbf{t}$. In this case, $k = 0$, and we must show that $\overline{\text{Val}}^{\leq 0}(v) \leq 0 \times W = 0$, which is true by definition of $\overline{\text{Val}}^{\leq 0}$. Or $v \neq \mathbf{t}$. In this case, $k > 0$, and we must show that $\overline{\text{Val}}^{\leq 0}(v) = +\infty$, which is true again by definition of $\overline{\text{Val}}^{\leq 0}$.

Inductive case: $j = \ell \geq 1$. Let us assume that the lemma holds for all v , for all values of j up to $\ell - 1$, and let us show that it holds for all v , and for $j = \ell$. Let us fix a vertex v , and its associated index k such that $v \in \text{Attr}_k(\{\mathbf{t}\}) \setminus \text{Attr}_{k-1}(\{\mathbf{t}\})$. We consider two cases.

1. First, assume $k > \ell$. In this case, we must show that $\overline{\text{Val}}^{\leq \ell}(v) = +\infty$. We consider again two cases:
 - (a) If $v \in V_{\text{Min}}$, then none of its successors belong to $\text{Attr}_{\ell-1}(\{\mathbf{t}\})$, otherwise, v would be in $\text{Attr}_{\ell}(\{\mathbf{t}\})$, by definition of the attractor, and we would have $k \leq \ell$. Hence, by induction hypothesis, $\overline{\text{Val}}^{\leq \ell-1}(v') = +\infty$ for all v' such that $(v, v') \in E$. Thus:

$$\begin{aligned} \overline{\text{Val}}^{\leq \ell}(v) &= \min_{(v, v') \in E} \left(\omega(v, v') + \overline{\text{Val}}^{\leq \ell-1}(v') \right) & (\text{Lemma 6}) \\ &= +\infty. \end{aligned}$$

- (b) If $v \in V_{\text{Max}}$, then at least one successor of v does not belong to $\text{Attr}_{\ell-1}(\{\mathbf{t}\})$, otherwise, v would be in $\text{Attr}_{\ell}(\{\mathbf{t}\})$, by definition of the attractor, and we would have $k \leq \ell$. Hence, by induction hypothesis, there exists v' such that $(v, v') \in E$ and $\overline{\text{Val}}^{\leq \ell-1}(v') = +\infty$. Thus:

$$\begin{aligned} \overline{\text{Val}}^{\leq \ell}(v) &= \max_{(v, v') \in E} \left(\omega(v, v') + \overline{\text{Val}}^{\leq \ell-1}(v') \right) & (\text{Lemma 6}) \\ &= +\infty. \end{aligned}$$

2. Second, assume $k \leq \ell$. In this case, we must show that $\overline{\text{Val}}^{\leq \ell}(v) \leq \ell W$. As in the previous item, we consider two cases:

- (a) In the case where $v \in V_{\text{Min}}$, we let $\bar{v}$ be a vertex such that $\bar{v} \in \text{Attr}_{k-1}(\{\mathbf{t}\})$ and $(v, \bar{v}) \in E$. Such a vertex exists by definition of the attractor. By induction hypothesis, $\overline{\text{Val}}^{\leq \ell-1}(\bar{v}) \leq \ell W$. Then:

$$\begin{aligned} \overline{\text{Val}}^{\leq \ell}(v) &= \min_{(v, v') \in E} \left(\omega(v, v') + \overline{\text{Val}}^{\leq \ell-1}(v') \right) && \text{(Lemma 6)} \\ &\leq \omega(v, \bar{v}) + \overline{\text{Val}}^{\leq \ell-1}(\bar{v}) && ((v, \bar{v}) \in E) \\ &\leq \omega(v, \bar{v}) + (\ell - 1)W && \text{(Ind. Hyp.)} \\ &\leq W + (\ell - 1)W \\ &= \ell W. \end{aligned}$$

- (b) In the case where $v \in V_{\text{Max}}$, we know that all successors v' of v belong to $\text{Attr}_{k-1}(\{\mathbf{t}\})$ by definition of the attractor. By induction hypothesis, for all successors v' of v : $\overline{\text{Val}}^{\leq \ell-1}(v') \leq \ell W$. Hence:

$$\begin{aligned} \overline{\text{Val}}^{\leq \ell}(v) &= \max_{(v, v') \in E} \left(\omega(v, v') + \overline{\text{Val}}^{\leq \ell-1}(v') \right) && \text{(Lemma 6)} \\ &\leq \max_{(v, v') \in E} (W + (\ell - 1)W) && \text{(Ind. Hyp.)} \\ &= \ell W. \end{aligned} \quad \square$$

In particular, this allows us to conclude that, after $|V|$ steps, all values are bounded by $|V|W$:

Corollary 8 *For all $v \in V$, $\overline{\text{Val}}^{\leq |V|}(v) \leq |V|W$.*

The next step is to show that the sequence $(x_i)_{i \geq 0}$ stabilises after a bounded number of steps, when all values are finite:

Lemma 9 *In an MCR game where all values are finite, the sequence $(\overline{\text{Val}}^{\leq i})_{i \geq 0}$ stabilises after at most $(2|V| - 1)W|V| + |V|$ steps.*

Proof We first show that if Min can secure, from some vertex v , a payoff less than $-(|V| - 1)W$, i.e. $\text{Val}(v) < -(|V| - 1)W$, then it can secure an arbitrarily small payoff from that vertex, i.e. $\text{Val}(v) = -\infty$, which contradicts our hypothesis that the value is finite. Hence, let us suppose that there exists a strategy σ_{Min} for Min such that $\text{Val}(v, \sigma_{\text{Min}}) < -(|V| - 1)W$. Let $\mathcal{G}'$ be the mean-payoff game studied in Proposition 4. We will show that $\text{Val}_{\mathcal{G}'}(v) < 0$, which permits to conclude that $\text{Val}_{\mathcal{G}}(v) = -\infty$. Let σ_{Max} be a memoryless strategy of Max. By hypothesis, we know that $\text{MCR}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})) < -(|V| - 1)W$. This ensures the existence of a cycle with negative cost in the play $\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})$: otherwise, we could iteratively remove every possible non-negative cycle of the finite play before reaching $\mathbf{t}$ (hence reducing the cost of the play) and obtain a play without cycles before reaching $\mathbf{t}$ with a cost less than $-(|V| - 1)W$, which is impossible (since it should be of length at most $|V| - 1$ to cross at most one occurrence of each vertex). Consider the first negative cycle in the play. After the first occurrence of the cycle, we let Min choose its actions like in the cycle. By this way, we can construct another strategy σ'_{Min} for Min, verifying that for every memoryless strategy σ_{Max} of Max (this would not be true for general strategies of Max), we have $\text{MP}(\text{Play}(v, \sigma_{\text{Max}}, \sigma'_{\text{Min}}))$ being

the weight of the negative cycle in which the play finishes. Since for mean-payoff games, memoryless strategies are sufficient for **Max**, we deduce that $\text{Val}_{\mathcal{G}'}(v) < 0$.

This reasoning permits to prove that at every step i , $\overline{\text{Val}}^{\leq i}(v) \geq \text{Val}(v) \geq -(|V| - 1)W + 1$ for all vertices v . Recall from Corollary 8 that, after $|V|$ steps in the sequence, all vertices are assigned a value smaller than $|V|W$. Moreover, we know that the sequence is non-increasing (see (1)). In summary, for all $k \geq 0$ and for all vertices v :

$$-(|V| - 1)W + 1 \leq \overline{\text{Val}}^{\leq |V|+k}(v) \leq |V|W$$

Hence, in the worst case a strictly decreasing sequence will need $(2|V| - 1)W|V|$ steps to reach the lowest possible value where all vertices are assigned $-(|V| - 1)W + 1$ from the highest possible value where all vertices are assigned $|V|W$. Thus, taking into account the $|V|$ steps to reach a finite value on all vertices, the sequence stabilises in at most $(2|V| - 1)W|V| + |V|$ steps. $\square$

Let us thus denote by $\overline{\text{Val}}^{\leq}$ the value obtained when the sequence $(\overline{\text{Val}}^{\leq i})_{i \geq 0}$ stabilises. We are now ready to prove that this value is the actual value of the game:

Lemma 10 *For all MCR games where all values are finite: $\overline{\text{Val}}^{\leq} = \text{Val}$.*

Proof We already know that $\overline{\text{Val}}^{\leq} \succcurlyeq \text{Val}$. Let us show that $\overline{\text{Val}}^{\leq} \preccurlyeq \text{Val}$. Let $v \in V$ be a vertex. Since $\text{Val}(v)$ is a finite integer, there exists a strategy σ_{Min} for **Min** that realises this value, i.e. $\text{Val}(v) = \sup_{\sigma_{\text{Max}}} \text{MCR}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}}))$. Notice that this holds because the values are integers, inducing that the infimum in the definition of $\overline{\text{Val}}(v) = \text{Val}(v)$ is indeed reached.

Let us build a tree $A_{\sigma_{\text{Min}}}$ unfolding all possible plays from v against σ_{Min} . $A_{\sigma_{\text{Min}}}$ has a root labeled by v . If a tree node is labeled by a vertex v of **Min**, this tree node has a unique child labeled by $\sigma_{\text{Min}}(v)$. If a tree node is labeled by a vertex v of **Max**, this tree node has one child per successor v' of v in the graph, labeled by v' . We proceed this way until we encounter a node labeled by a vertex from $\mathbf{t}$ in which case this node is a leaf. $A_{\sigma_{\text{Min}}}$ is necessarily finite. Otherwise, by König's Lemma, it has one infinite branch that never reaches $\mathbf{t}$. From that infinite branch, one can extract a strategy σ_{Max} for **Max** such that $\text{MCR}(\text{Play}(v, \sigma_{\text{Max}}, \sigma_{\text{Min}})) = +\infty$, hence $\text{Val}(v) = +\infty$, which contradicts the hypothesis. Assume the tree has depth m . Then, $A_{\sigma_{\text{Min}}}$ is a subtree of the tree A obtained by unfolding all possible plays up to length m (as in the proof of Lemma 6). In this case, it is easy to check that the value labeling the root of $A_{\sigma_{\text{Min}}}$ after applying backward induction is larger than or equal to the value labeling the root of A after applying backward induction. The latter is $\text{Val}(v)$ while the former is $\overline{\text{Val}}^{\leq m}(v)$, by Lemma 6, so that $\text{Val}(v) \geq \overline{\text{Val}}^{\leq m}(v)$. Since the sequence is non-increasing, we finally obtain $\text{Val}(v) \geq \overline{\text{Val}}^{\leq}(v)$. $\square$

As a corollary of this lemma, we obtain:

Corollary 11 *In all MCR games where all values are finite, Val is the greatest fixed point of $\mathcal{F}$.*

We are finally able to establish the correctness of Algorithm 1.

Proof (of Proposition 5) Let us first suppose that the values of all vertices are finite. Then, $x_j = \overline{\text{Val}}^{\leq j}$ is the value of X at the beginning of the j th step of the loop, and the condition of line 13 can never be fulfilled. Hence, by Lemma 9, after at most $(2|V| - 1)W|V| + |V|$ iterations, all values are computed correctly (by Lemma 10) in that case.

Suppose now that there are vertices with value $+\infty$. Those vertices will remain at their initial value $+\infty$ during the whole computation, and hence do not interfere with the rest of the computation.

Finally, consider that the game contains vertices with value $-\infty$. By the proof of Lemma 9, we know that optimal values of vertices of values different from $-\infty$ are at least $-(|V| - 1)W + 1$ so that, if the value of a vertex reaches an integer below $-(|V| - 1)W$, we are sure that its value is indeed $-\infty$, which proves correct the line 13 of the algorithm. This update may cost at most one step per vertex, which in total adds at most $|V|$ iterations. Moreover, dropping the value to $-\infty$ does not harm the correction for the other vertices (it may only speed up the convergence of their values). This is due to the fact that, if the Kleene sequence $(\mathcal{F}^i(x_0))_{i \geq 0}$ is initiated with a vector of values x_0 that is greater or equal to the optimal value vector Val , then the sequence converges at least as fast as before towards the optimal value vector. $\square$

Example 12 We close this discussion on the computation of the values by an example of execution of Algorithm 1. Consider the MCR game in Fig. 2. The successive values for vertices (v_1, v_2) (value of the target v_3 is always 0) computed by the value iteration algorithm are the following: $(+\infty, +\infty)$, $(+\infty, 0)$, $(-1, 0)$, $(-1, -1)$, $(-2, -1)$, $(-2, -2)$, $\dots$, $(-W, -W + 1)$, $(-W, -W)$. This requires $2W$ steps to converge (hence a pseudo-polynomial time).

3.4 Computing optimal strategies for both players

As we have seen earlier, Min does not always have optimal memoryless strategies. However, we will see that one can always construct so-called *negative cycle strategies* (NC-strategies), which are memoryless strategies that have a meaningful structure for Min, in the sense that they allow him either: (i) to reach the target by means of a play whose value is lower than the value of the game; or (ii) to decrease arbitrarily the partial sums along the play, when it does not reach the target (in other words, the partial sums tend to $-\infty$ as the play goes on). So, NC-strategies are, in general, not optimal, as they do not guarantee to reach the target (but in this case, they guarantee that Min will play consistently with his objective, by decreasing the value of the play prefixes).

Formally, an NC-strategy is a memoryless strategy σ_{Min} for player Min such that, for all cycles $v_0 \dots v_{k-1} v_k$ (with $v_k = v_0$) conforming to σ_{Min} (i.e. where each $v_i \in V_{\text{Min}}$ satisfies $\sigma(v_i) = v_{i+1}$),

$$\omega(v_0, v_1) + \dots + \omega(v_{k-1}, v_k) < 0.$$

We define the *fake value*, $\text{fake}(v, \sigma_{\text{Min}})$, of an NC-strategy σ_{Min} as the supremum of the costs of the *finite* plays that conform with it and start in v :

$$\text{fake}(v, \sigma_{\text{Min}}) = \sup \{ \mathbf{MCR}(v_1 \dots v_k) \mid v_1 = v, v_k = \mathbf{t}, \forall v_i \in V_{\text{Min}} : \sigma_{\text{Min}}(v_i) = v_{i+1} \}.$$

Notice that the fake value is not necessarily equal to the value of σ_{Min} , since, in the definition of the fake value, we only consider plays that *do reach* the target, and ignore those that don't (in the computation of the actual value of σ_{Min} , these plays would yield $+\infty$). However, a strategy's fake value is always smaller than or equal to its value. We say that an NC-strategy is *fake-optimal* if its fake value is smaller than or equal to the optimal value of the game from all vertices. In particular, if the optimal value of a vertex v is $-\infty$, the set $\{\mathbf{MCR}(v_1 \cdots v_k) \mid v_1 = v, v_k = \mathbf{t}, \forall v_i \in V_{\text{Min}} \sigma_{\text{Min}}(v_i) = v_{i+1}\}$ is empty for all fake-optimal strategies σ_{Min} , hence the supremum of this set is indeed $-\infty$.

Example 13 On the game of Fig. 2, the memoryless strategy σ_{Min} mapping v_2 to v_1 is an NC-strategy as the only two possible cycles $v_1 v_2 v_1$ and $v_2 v_1 v_2$ both have weight -1 . The value of σ_{Min} from v_2 is $+\infty$ as the play $v_2 v_1 v_2 v_1 v_2 \cdots$ agrees with it and does not reach the target, but the fake value of σ_{Min} is $-W$, since the play $v_2 v_1 v_3$ is the finite play that agrees with σ_{Min} with the biggest cost possible. Since we know that the actual optimal value is $-W$, the strategy σ_{Min} is fake-optimal. $\square$

The following proposition reveals the interest of NC-strategies, by explaining how one can, in some cases, construct an optimal finite-memory strategy from a fake-optimal NC-strategy.

Proposition 14 *If Min has a strategy $\sigma_{\text{Min}}^\dagger$ to reach a target vertex (from every possible initial vertex), and has a fake-optimal NC-strategy σ_{Min}^* , then for all $n \in \mathbb{Z}$ one can construct a finite-memory strategy σ_{Min}^n such that for all vertices v , it holds that $\text{Val}(v, \sigma_{\text{Min}}^n) \leq \max(n, \text{Val}(v))$.*

Remark 15 In particular, if the value of all vertices is finite, then one can construct an optimal finite-memory strategy. If the value of a vertex is $-\infty$, this proposition also says that there is an infinite family of strategies that allows one to secure a value which is arbitrarily low (remember that, by definition, $-\infty$ can not be the value that corresponds to a single strategy).

Proof First let us show that for all partial plays $\pi = v_1 \cdots v_\ell$ of size at least $k|V| + 1$ (for some k) that conforms with σ_{Min}^* , $\mathbf{TP}(\pi) \leq W(|V| - 1) - k$. We establish this proof by induction on k . For the base case, we consider the case where $\ell \leq |V|$. Then, the play visits at most $|V| - 1$ edges, and thus its total cost is at most $W(|V| - 1)$. Now, for the induction, we assume that $\ell \geq k|V| + 1$ for some $k \geq 1$. Then, let i and j be two indices such that $i < j$, $v_i = v_j$ and $j \leq i + |V|$ (those indices necessarily exist since $\ell \geq |V| + 1$). Since σ_{Min}^* is an NC-strategy, the total cost of $v_i \cdots v_j$ is at most -1 . As $\pi' = v_1 \cdots v_i v_{j+1} \cdots v_\ell$ is also a play that conforms with σ_{Min}^* with size greater than $(k-1)|V| + 1$, we have (by induction hypothesis) $\mathbf{TP}(\pi') \leq W(|V| - 1) - k + 1$, thus $\mathbf{TP}(\pi) = \mathbf{TP}(v_i \cdots v_j) + \mathbf{TP}(\pi') \leq W(|V| - 1) - k$.

In the following, let $\sigma_{\text{Min}}^\dagger$ be a memoryless strategy ensuring to reach the target (obtained by the attractor technique for instance), and let $k = \max(2W(|V| - 1) - n, 0)$. The strategy σ_{Min}^n consists in playing σ_{Min}^* , until switching to $\sigma_{\text{Min}}^\dagger$ when the length of the play is greater than $k|V| + 1$: formally $\sigma_{\text{Min}}^n(\pi) = \sigma_{\text{Min}}^*(\pi)$ if $|\pi| < k|V| + 1$ and $\sigma_{\text{Min}}^n(\pi) = \sigma_{\text{Min}}^\dagger(\pi)$ otherwise. It is clear that this strategy can be implemented by a finite deterministic Moore machine, storing the size of the current play until it is greater than $k|V| + 1$.

Let us now check that $\text{Val}(v, \sigma_{\text{Min}}^n) \leq \max(n, \text{Val}(v))$ for all vertices. Let π be a play conforming to σ_{Min}^n starting in a vertex v and reaching $\mathbf{t}$. If $|\pi| \leq k|V| + 1$ then π is a play that reaches the target conforming to σ_{Min}^* and therefore

$$\begin{aligned} \text{MCR}(\pi) &\leq \text{fake}(v, \sigma_{\text{Min}}^*) && \pi \text{ is a finite play reaching } \mathbf{t} \\ &\leq \text{Val}(v) && \sigma_{\text{Min}}^* \text{ is fake-optimal} \\ &\leq \max(n, \text{Val}(v)). \end{aligned}$$

If $|\pi| > k|V| + 1$, then let π_1 be its prefix of size $k|V| + 1$, and π_2 be the rest of the play ($\pi = \pi_1 \cdot \pi_2$). As π_2 is a play that conforms with $\sigma_{\text{Min}}^\dagger$, a memoryless strategy ensuring to reach the target, we know that it reaches the target in at most $|V|$ steps. Therefore, $\text{TP}(\pi_2) \leq W(|V| - 1)$. As π_1 is a play that conforms with σ_{Min}^* of size $k|V| + 1$, from the reasoning above, $\text{TP}(\pi_1) \leq W(|V| - 1) - k \leq n - W(|V| - 1)$, therefore $\text{TP}(\pi) = \text{TP}(\pi_1) + \text{TP}(\pi_2) \leq n \leq \max(n, \text{Val}(v))$. $\square$

In practice, rather than using a Moore machine, we can simulate the strategy σ_{Min}^n (of the proof above) quite easily: one just has to handle two memoryless strategies and a counter keeping track of the length of the current play. Since $\sigma_{\text{Min}}^\dagger$ can easily be obtained by the classical attractor algorithm, we turn our attention to the construction of a fake-optimal NC-strategy σ_{Min}^* . Without loss of generality, we suppose that no vertex has optimal value $+\infty$, since for these vertices, all strategies are equivalent.

For vertices of value $-\infty$, we can obtain σ_{Min}^* as an optimal strategy for Min in the mean-payoff game of the first item in Proposition 4. Since the mean-payoff value is negative, this strategy guarantees that it does not reach target, thus generating a fake value $-\infty$, equal to the optimal value of the vertex. Moreover, since it is a memoryless strategy, we know that, as soon as Max plays a memoryless strategy that necessarily reaches a cycle, this cycle must have a negative weight (at most the optimal value of the initial vertex) : this strategy is thus a fake-optimal NC-strategy.

From now on, we thus concentrate our study on the vertices of finite value, thus considering that no vertices have value $+\infty$ or $-\infty$ in the MCR game. Let X^i denote the value of variable X after i iterations of the loop of Algorithm 1, and let $X^0(v) = +\infty$ for all $v \in V$. We have seen that the sequence $X^0 \succcurlyeq X^1 \succcurlyeq X^2 \succcurlyeq \dots$ is stationary at some point, equal to Val . Let us now define $\sigma_{\text{Min}}^*(v)$ for all vertices $v \in V_{\text{Min}} \setminus \{\mathbf{t}\}$. We let $i_v > 0$ be the smallest index such that $X^{i_v}(v) = \text{Val}(v)$. Fix a vertex v' such that $X^{i_v}(v) = \omega(v, v') + X^{i_v-1}(v')$ (such a v' exists by definition) and define $\sigma_{\text{Min}}^*(v) = v'$. Notice that it is exactly what is achieved in line 11 of Algorithm 1. Note also that strategy $\sigma_{\text{Min}}^\dagger$ is correctly computed in line 12 of Algorithm 1.

Let us prove that this construction indeed yields a fake-optimal NC strategy σ_{Min}^* . We first prove the following lemma, that states that the vertex $\sigma_{\text{Min}}^*(v)$ has already reached its final value at step $i_v - 1$ of the algorithm, for all vertices v .

Lemma 16 *For all vertices $v \in V_{\text{Min}} \setminus \{\mathbf{t}\}$, $X^{i_v-1}(\sigma_{\text{Min}}^*(v)) = \text{Val}(\sigma_{\text{Min}}^*(v))$.*

Proof Let $v' = \sigma_{\text{Min}}^*(v)$. Since $(X^i(v'))_{i \geq 0}$ is non-increasing and converges towards the value of v' , we know that $X^{i_v-1}(v') \geq \text{Val}(v')$. By contradiction assume that

$X^{i_v-1}(v') > \text{Val}(v')$. Note that there exists $j > i_v$ such that $X^{j-1}(v') = \text{Val}(v')$. Therefore:

$$\begin{aligned}
\text{Val}(v) &\leq X^j(v) && (\text{since } (X^k(v))_{k \in \mathbb{N}} \text{ is non-increasing}) \\
&\leq \omega(v, v') + X^{j-1}(v') && (\text{by definition of } X^j) \\
&= \omega(v, v') + \text{Val}(v') && (\text{by choice of } j) \\
&< \omega(v, v') + X^{i_v-1}(v') && (\text{by the contradiction hypothesis}) \\
&= X^{i_v}(v) = \text{Val}(v) && (\text{by definition of } X^{i_v})
\end{aligned}$$

which raises a contradiction. $\square$

We can now prove that our definition of σ_{Min}^* has the announced properties:

Proposition 17 σ_{Min}^* is an NC-strategy, and $\text{fake}(v, \sigma_{\text{Min}}^*) \leq \text{Val}(v)$ for all vertices v .

Proof Equivalently, we want to prove that for all vertices v , and for all plays $\pi = v_1 v_2 \dots$ starting in v and conforming to σ_{Min}^* ,

- (1) if there exists $i < j$ such that $v_i = v_j$, then $\mathbf{TP}(v_i \dots v_j) < 0$,
- (2) if π reaches $\mathfrak{t}$ then $\mathbf{MCR}(\pi) \leq \text{Val}(v)$.

For (1), consider a cycle $v_i \dots v_j$ with $v_j = v_i$. Notice that at least one vertex of this cycle belongs to Min , since, otherwise, Max would have a strategy to obtain a value $+\infty$ for vertex v_i , which contradicts the hypothesis that every vertex has value different from $+\infty$. Hence, for the sake of the explanation, we suppose that $v_i \in V_{\text{Min}}$, and that the index i_{v_i} where vertex v_i stabilises in the sequence $(X^k(v_i))_{k \geq 0}$ is maximal among all possible vertices of $\{v_i, \dots, v_j\} \cap V_{\text{Min}}$. The following extends easily to the general case.

We prove by induction over $i < \ell \leq j$ that

$$X^{i_{v_i}}(v_i) \geq \mathbf{TP}(v_i \dots v_\ell) + X^{i_{v_i}-1}(v_\ell). \quad (3)$$

The base case $\ell = i + 1$ comes from the fact that, since $v_i \in V_{\text{Min}}$, we have $\sigma_{\text{Min}}^*(v_i) = v_{i+1}$. Therefore, $X^{i_{v_i}}(v_i) = \omega(v_i, v_{i+1}) + X^{i_{v_i}-1}(v_{i+1})$, and by definition $\mathbf{TP}(v_i v_{i+1}) = \omega(v_i, v_{i+1})$.

For the inductive case, let us consider $i < \ell < j$ such that (3) holds and let us prove it for $\ell + 1$. If $v_\ell \in V_{\text{Max}}$, by definition of $X^{i_{v_i}}$, we have

$$X^{i_{v_i}}(v_\ell) = \max_{(v_\ell, v') \in E} \omega(v_\ell, v') + X^{i_{v_i}-1}(v') \geq \omega(v_\ell, v_{\ell+1}) + X^{i_{v_i}-1}(v_{\ell+1}).$$

If $v_\ell \in V_{\text{Min}}$, by maximality of i_{v_i} , we have

$$\begin{aligned}
X^{i_{v_i}}(v_\ell) &= X^{i_{v_\ell}}(v_\ell) = \omega(v_\ell, v_{\ell+1}) + X^{i_{v_\ell}-1}(v_{\ell+1}) \\
&\geq \omega(v_\ell, v_{\ell+1}) + X^{i_{v_i}-1}(v_{\ell+1})
\end{aligned}$$

using that the sequence $X^0, X^1, X^2, \dots$ is non-increasing. In all cases, we have

$$X^{i_{v_i}}(v_\ell) \geq \omega(v_\ell, v_{\ell+1}) + X^{i_{v_i}-1}(v_{\ell+1}).$$

Using again that $X^0, X^1, X^2, \dots$ is non-increasing, we obtain

$$X^{i_{v_i}-1}(v_\ell) \geq X^{i_{v_i}}(v_\ell) \geq \omega(v_\ell, v_{\ell+1}) + X^{i_{v_i}-1}(v_{\ell+1}).$$

Injecting this into the induction hypothesis, we have

$$\begin{aligned} X^{i_{v_i}}(v_i) &\geq \mathbf{TP}(v_i \cdots v_\ell) + \omega(v_\ell, v_{\ell+1}) + X^{i_{v_i}-1}(v_{\ell+1}) \\ &= \mathbf{TP}(v_i \cdots v_{\ell+1}) + X^{i_{v_i}-1}(v_{\ell+1}) \end{aligned}$$

which concludes the proof by induction.

In particular, for $\ell = j$, as $v_i = v_j$ we obtain that

$$X^{i_{v_i}}(v_i) \geq X^{i_{v_i}-1}(v_i) + \mathbf{TP}(v_i \cdots v_j).$$

and as, by definition of i_{v_i} , we have $X^{i_{v_i}}(v_i) < X^{i_{v_i}-1}(v_i)$, we necessarily have $\mathbf{TP}(v_i \cdots v_j) < 0$.

We now prove (2). We decompose π as $\pi = v_1 \cdots v_k \mathbf{t}^\omega$ with for all i , $v_i \neq \mathbf{t}$. We prove by decreasing induction on i that $\mathbf{MCR}(v_i \cdots v_k \mathbf{t}^\omega) \leq \mathbf{Val}(v_i)$. If $i = k + 1$, $\mathbf{MCR}(\mathbf{t}^\omega) = 0 = \mathbf{Val}(\mathbf{t})$. If $i \leq k$, by induction we have $\mathbf{MCR}(v_{i+1} \cdots v_k \mathbf{t}^\omega) \leq \mathbf{Val}(v_{i+1})$. Thus $\mathbf{MCR}(v_i \cdots v_k \mathbf{t}^\omega) \leq \omega(v_i, v_{i+1}) + \mathbf{Val}(v_{i+1})$. If $v_i \in V_{\text{Min}}$, then $v_{i+1} = \sigma_{\text{Min}}^*(v_i)$, and by Lemma 16,

$$\mathbf{Val}(v_{i+1}) = X^{i_{v_i}-1}(v_{i+1}) = X^{i_{v_i}}(v_i) - \omega(v_i, v_{i+1})$$

by definition of σ_{Min}^* . Since $X^{i_{v_i}}(v_i) = \mathbf{Val}(v_i)$, this can be rewritten as $\mathbf{Val}(v_i) = \omega(v_i, v_{i+1}) + \mathbf{Val}(v_{i+1}) \geq \mathbf{MCR}(v_i \cdots v_k \mathbf{t}^\omega)$, which proves the inequality for i . If $v_i \in V_{\text{Max}}$, then

$$\mathbf{MCR}(v_i \cdots v_k \mathbf{t}^\omega) \leq \omega(v_i, v_{i+1}) + \mathbf{Val}(v_{i+1}) \leq \max_{(v_i, v') \in E} \omega(v_i, v') + \mathbf{Val}(v') = \mathbf{Val}(v_i)$$

the last equality coming from Corollary 11. $\square$

As a corollary, we obtain the existence of *finite memory strategies* (obtained from σ_{Min}^* and $\sigma_{\text{Min}}^\dagger$ as described above) when all values are finite:

Corollary 18 *When the values of all vertices of an MCR game are finite, one can construct an optimal finite-memory strategy for player Min.*

Strategies of Max Let us now show that Max always has a *memoryless optimal strategy*. This asymmetry stems directly from the asymmetric definition of the game—while Min has the double objective of reaching $\mathbf{t}$ and minimising its cost, Max aims at avoiding $\mathbf{t}$, and if not possible, maximising the cost.

Proposition 19 *In all MCR games, Max has a memoryless optimal strategy.*

Proof For vertices with value $+\infty$, we already know a memoryless optimal strategy for Max, namely every strategy that remains outside the attractor of the target vertices. For vertices with value $-\infty$, all strategies are equally bad for Max.

We now explain how to define a memoryless optimal strategy σ_{Max}^* for Max in case of a graph containing only finite values. For every finite play π ending in a vertex $v \in V_{\text{Max}}$ of Max, we let

$$\sigma_{\text{Max}}^*(\pi) = \operatorname{argmax}_{v' \in E(v)} (\omega(v, v') + \mathbf{Val}(v')).$$

This is clearly a memoryless strategy. Let us prove that it is optimal for Max, that is, for every vertex $v \in V$, and every strategy σ_{Min} of Min

$$\mathbf{MCR}(\text{Play}(v, \sigma_{\text{Max}}^*, \sigma_{\text{Min}})) \geq \mathbf{Val}(v).$$

In case $\text{Play}(v, \sigma_{\text{Max}}^*, \sigma_{\text{Min}})$ does not reach the target set of vertices, the inequality holds trivially. Otherwise, we let $\text{Play}(v, \sigma_{\text{Max}}^*, \sigma_{\text{Min}}) = v_0 v_1 \cdots v_\ell \cdots$ with ℓ the least position such that $v_\ell = \mathbf{t}$. If $\ell = 0$, i.e. $v = v_0 = \mathbf{t}$, we have

$$\mathbf{MCR}(\text{Play}(v, \sigma_{\text{Max}}^*, \sigma_{\text{Min}})) = 0 = \mathbf{Val}(v).$$

Otherwise, let us prove by induction on $0 \leq i \leq \ell$ that

$$\mathbf{MCR}(v_{\ell-i} \cdots v_\ell) \geq \mathbf{Val}(v_{\ell-i}).$$

This will permit to conclude since

$$\mathbf{MCR}(\text{Play}(v, \sigma_{\text{Max}}^*, \sigma_{\text{Min}})) = \mathbf{MCR}(v_0 v_1 \cdots v_\ell) \geq \mathbf{Val}(v_0) = \mathbf{Val}(v).$$

The base case $i = 0$ corresponds to the previous case where the starting vertex is $\mathbf{t}$. Supposing that the property holds for index i , let us prove it for $i + 1$. We have

$$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) = \omega(v_{\ell-i-1}, v_{\ell-i}) + \mathbf{MCR}(v_{\ell-i} \cdots v_\ell).$$

By induction hypothesis, we have

$$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) \geq \omega(v_{\ell-i-1}, v_{\ell-i}) + \mathbf{Val}(v_{\ell-i}). \quad (4)$$

We now consider two cases:

- If $v_{\ell-i-1} \in V_{\text{Max}} \setminus \{\mathbf{t}\}$, then $v_{\ell-i} = \sigma_{\text{Max}}^*(v_0 v_1 \cdots v_{\ell-i-1})$, so that by definition of σ_{Max}^* :

$$\omega(v_{\ell-i-1}, v_{\ell-i}) + \mathbf{Val}(v_{\ell-i}) = \max_{v' \in V \mid (v_{\ell-i-1}, v') \in E} (\omega(v_{\ell-i-1}, v') + \mathbf{Val}(v')).$$

Using Corollary 11 and (4), we obtain

$$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) \geq \mathbf{Val}(v_{\ell-i-1}).$$

- If $v_{\ell-i-1} \in V_{\text{Min}} \setminus \{\mathbf{t}\}$, then

$$\omega(v_{\ell-i-1}, v_{\ell-i}) + \mathbf{Val}(v_{\ell-i}) \geq \min_{v' \in V \mid (v_{\ell-i-1}, v') \in E} (\omega(v_{\ell-i-1}, v') + \mathbf{Val}(v')).$$

Once again using Corollary 11 and (4), we obtain

$$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) \geq \mathbf{Val}(v_{\ell-i-1}).$$

This concludes the proof. $\square$

Fig. 3 Importance of the **if** condition of Algorithm 1

Notice that strategy σ_{Max}^* can be computed directly, along the execution of the value iteration algorithm. This corresponds to line 7 of Algorithm 1.

Notice further that the presence of the **if** condition in line 11, and the absence of a similar condition at line 7, are crucial. Indeed, removing the **if** from line 11 would amount to computing σ_{Min}^* from the vector of values obtained at the end of the value iteration, when the vector $\mathbf{X}$ has stabilised. Let us show that, in this case, the algorithm might fail to compute a fake-optimal NC-strategy, by considering the MCR game in the left part of Fig. 3. Clearly, the values of both v_1 and v_2 are 0. However, if we extract σ_{Min}^* from $\mathbf{X}_{\text{pre}}$ at that point of the execution of the algorithm (i.e. we let $\sigma_{\text{Min}}^*(v) = \text{argmin}_{v' \in E(v)} (\omega(v, v') + \mathbf{X}_{\text{pre}}(v'))$), then we could end up with $\sigma_{\text{Min}}^*(v_1) = v_2$ and $\sigma_{\text{Min}}^*(v_2) = v_1$. In this case, σ_{Min}^* is no longer an NC-strategy because the cycle $v_1 v_2 v_1$ does not have negative weight. Similarly, let us consider the MCR game in the right part of Fig. 3 to explain why line 7 is not under the range of an **if**($\mathbf{X}(v) \neq \mathbf{X}_{\text{pre}}(v)$) condition. After two iterations, $\mathbf{X}$ reaches the optimal values $(-1, 0, -1, 0)$ but a Max strategy σ_{Max}^* such that $\sigma_{\text{Max}}^*(v_1) = v_3$ can still be chosen since $\mathbf{X}_{\text{pre}}(v_3) = 0$ at that point. However, on the next iteration, $\mathbf{X}(v_3) = \mathbf{X}_{\text{pre}}(v_3) = -1$ (indeed, $\mathbf{X}$ has now stabilised on all nodes), and it is crucial that $\sigma_{\text{Max}}^*(v_1) = v_2$ gets computed, otherwise the strategy would not be optimal for Max.

4 An efficient algorithm to solve total-payoff games

We now turn our attention back to total-payoff games (without reachability objective), and discuss our main contribution. Building on the results of the previous section, we introduce the *first* (as far as we know) *pseudo-polynomial time algorithm* for solving those games in the presence of arbitrary weights, thanks to a reduction from total-payoff games to MCR games. The MCR game produced by the reduction has size pseudo-polynomial in the size of the original total-payoff game. Then, we show how to compute the values of the total-payoff game without building the entire MCR game, and explain how to deduce memoryless optimal strategies from the computation of our algorithm.

4.1 Reduction to MCR games

We provide a transformation from a total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP} \rangle$ to an MCR game $\mathcal{G}^K$ (where K is a parameter in $\mathbb{N}$) such that the values of $\mathcal{G}$ can be extracted from the values in $\mathcal{G}^K$ (as formalised below). Intuitively, $\mathcal{G}^K$ simulates the game where players play in $\mathcal{G}$; Min may propose to stop playing and reach a

Fig. 4 MCR game $\mathcal{G}^3$ associated with the total-payoff game of Fig. 2

fresh vertex $\mathbf{t}$ acting as the target; Max can then accept, in which case we reach the target, or refuse at most K times, in which case the game continues. Structurally, $\mathcal{G}^K$ consists of a sequence of copies of $\mathcal{G}$ along with some new states that we now describe formally. We let $\mathbf{t}$ be a fresh vertex, and, for all $n \geq 1$, we define the MCR game $\mathcal{G}^n = \langle V^n, E^n, \omega^n, \{\mathbf{t}\}\text{-MCR} \rangle$ where V_{Max}^n (respectively, V_{Min}^n) consists of n copies (v, j) , with $1 \leq j \leq n$, of each vertex $v \in V_{\text{Max}}$ (respectively, $v \in V_{\text{Min}}$) and some *exterior vertices* $(\mathbf{ex}, v, j)$ for all $v \in V$ and $1 \leq j \leq n$ (respectively, *interior vertices* $(\mathbf{in}, v, j)$ for all $v \in V$ and $1 \leq j \leq n$). Moreover, V_{Max}^n contains the fresh target vertex $\mathbf{t}$. Edges are given by

$$\begin{aligned} E^n = & \{(\mathbf{t}, \mathbf{t})\} \uplus \{((v, j), (\mathbf{in}, v', j)) \mid (v, v') \in E, 1 \leq j \leq n\} \\ & \uplus \{((\mathbf{in}, v, j), (v, j)) \mid v \in V, 1 \leq j \leq n\} \uplus \{((\mathbf{ex}, v, j), \mathbf{t}) \mid v \in V, 1 \leq j \leq n\} \\ & \uplus \{((\mathbf{in}, v, j), (\mathbf{ex}, v, j)) \mid v \in V, 1 \leq j \leq n\} \\ & \uplus \{((\mathbf{ex}, v, j), (v, j-1)) \mid v \in V, 1 < j \leq n\}. \end{aligned}$$

All edge weights are zero, except edges $((v, j), (\mathbf{in}, v', j))$ that have weight $\omega(v, v')$.

Example 20 For example, considering the weighted graph of Fig. 2, the corresponding reachability total-payoff game $\mathcal{G}^3$ is depicted in Fig. 4 (where weights 0 have been removed).

The next proposition formalises the relationship between the two games, and is proved in the rest of this subsection.

Proposition 21 *Let $K = |V|(2(|V| - 1)W + 1)$. For all $v \in V$ and $k \geq K$,*

- $\text{Val}_{\mathcal{G}}(v) \neq +\infty$ if and only if $\text{Val}_{\mathcal{G}}(v) = \text{Val}_{\mathcal{G}^k}((v, k))$;
- $\text{Val}_{\mathcal{G}}(v) = +\infty$ if and only if $\text{Val}_{\mathcal{G}^k}((v, k)) \geq (|V| - 1)W + 1$.

The bound K will be found by using the fact (informally described in the previous section) that if not infinite, the value of a min-cost reachability game belongs in $[-(|V| - 1) \times W + 1, |V| \times W]$, and that after enough visits of the same vertex, an adequate loop ensures that $\mathcal{G}^k$ verifies the above properties.

To prove Proposition 21, we must relate paths in games $\mathcal{G}$ and $\mathcal{G}^n$: with each finite path in $\mathcal{G}^n$, we associate a finite path in $\mathcal{G}$, obtained by looking at the

sequence of vertices of V appearing inside the vertices of the finite play. Formally, the *projection* of a finite path π is the sequence $\text{proj}(\pi)$ of vertices of $\mathcal{G}$ inductively defined by $\text{proj}(\varepsilon) = \varepsilon$ and for all finite path π , $v \in V$ and $1 \leq j \leq n$:

$$\begin{aligned} \text{proj}((\text{in}, v, j)\pi) &= \text{proj}(\pi), & \text{proj}((\text{ex}, v, j)\mathbf{t}\pi) &= v, & \text{proj}((\text{ex}, v, j)) &= \varepsilon, \\ \text{proj}((\text{ex}, v, j+1)(v, j)\pi) &= \text{proj}((v, j)\pi) = v \text{proj}(\pi). \end{aligned}$$

In particular, notice that in the case of a play with prefix $(\text{ex}, v, j)\mathbf{t}$, the rest of the play is entirely composed of target vertices $\mathbf{t}$, since $\mathbf{t}$ is a sink state. For instance, the projection of the finite play

$$(v_1, 3)(\text{in}, v_2, 3)(\text{ex}, v_2, 3)(v_2, 2)(\text{in}, v_3, 2)(v_3, 2)(\text{in}, v_3, 2)(\text{ex}, v_3, 2)\mathbf{t}$$

of the game $\mathcal{G}^3$ of Fig. 4 is given by $v_1 v_2 v_3 v_3$.

The following lemma relates plays of $\mathcal{G}^n$ with their projection in $\mathcal{G}$, comparing their total-payoff.

Lemma 22 *The projection mapping satisfies the following properties.*

1. If π is a finite play in $\mathcal{G}^n$ then $\text{proj}(\pi)$ is a finite play in $\mathcal{G}$.
2. If π is a play in $\mathcal{G}^n$ that does not reach the target, then $\text{proj}(\pi)$ is a play in $\mathcal{G}$.
3. For all finite play π , $\mathbf{TP}(\pi) = \mathbf{TP}(\text{proj}(\pi))$.

Proof The proof of 2 is a direct consequence of 1. With each vertex $w \in V^n \setminus \{\mathbf{t}\}$, we naturally associate a vertex $f(w)$ as follows:

$$f(v, j) = f(\text{in}, v, j) = f(\text{ex}, v, j) = v.$$

Then notice that if $(w, w') \in E^n$ with $w, w' \neq \mathbf{t}$, then either $f(w) = f(w')$ or $(f(w), f(w')) \in E$. We now prove 1 and 3 inductively on the size of the finite play $\pi = w_1 \cdots w_k$ of $\mathcal{G}^n$, along with the fact that

4. if $\text{proj}(\pi) \neq \emptyset$ and $w_1 \neq \mathbf{t}$ then the first vertex of $\text{proj}(\pi)$ is $f(w_1)$.

If $k = 0$, then $\pi = \text{proj}(\pi) = \varepsilon$ are finite plays with the same total-payoff. If $k = 1$, either $\text{proj}(\pi) = \varepsilon$ or $\pi = (v, j)$ and $\text{proj}(\pi) = v$: in both cases, the properties hold trivially. Otherwise, $k \geq 2$ and we distinguish several possible prefixes:

- If $\pi = (\text{in}, v, j)\pi'$, then $\text{proj}(\pi) = \text{proj}(\pi')$. Hence, 1 holds by induction hypothesis. If $\text{proj}(\pi)$ is non-empty, so is $\text{proj}(\pi')$. Moreover, the first vertex of π' is either (v, j) or (ex, v, j) , so that we have 4 by induction hypothesis. Finally, the previous remark shows that the first edge of π has necessarily weight 0, so that, $\mathbf{TP}(\pi) = \mathbf{TP}(\pi')$, and 3 also holds by induction hypothesis.
- If $\pi = (v, j)\pi'$, then $\text{proj}(\pi) = v \text{proj}(\pi')$ so that 4 holds directly. Moreover, π' is a non-empty finite play so that $\pi' = (\text{in}, v', j)\pi''$ with $(v, v') \in E$, and $\text{proj}(\pi') = \text{proj}(\pi'')$. By induction, $\text{proj}(\pi')$ is a finite play in $\mathcal{G}$, and it starts with v' (by 4). Since $(v, v') \in E$, this shows that $\text{proj}(\pi)$ is a finite play. Moreover, $\mathbf{TP}(\pi) = \omega^n((v, j), (\text{in}, v', j)) + \mathbf{TP}(\pi') = \omega(v, v') + \mathbf{TP}(\pi')$. By induction hypothesis, we have $\mathbf{TP}(\pi') = \mathbf{MCR}(\text{proj}(\pi'))$. Moreover, $\mathbf{MCR}(\text{proj}(\pi)) = \omega(v, v') + \mathbf{MCR}(\text{proj}(\pi'))$ which concludes the proof of 3.
- If $\pi = (\text{ex}, v, j)(v, j-1)\pi'$ then $\text{proj}(\pi) = v \text{proj}(\pi') = \text{proj}((v, j-1)\pi')$: this allows us to conclude directly by using the previous case.

- Otherwise, $\pi = (\mathbf{ex}, v, j)\mathbf{t}\pi'$, and then $\text{proj}(\pi) = v$ is a finite play with total-payoff 0, like π , and 4 holds trivially. $\square$

The next lemma states that when playing memoryless strategies, one can bound the total-payoff of all finite plays.

Lemma 23 *Let $v \in V$, and σ_{Min} (respectively, σ_{Max}) be a memoryless strategy for Min (respectively, Max) in the total-payoff game $\mathcal{G}$, such that $\text{Val}(v, \sigma_{\text{Min}}) \neq +\infty$ (respectively, $\text{Val}(v, \sigma_{\text{Max}}) \neq -\infty$). Then for all finite play π conforming to σ_{Min} (respectively, to σ_{Max}), $\text{TP}(\pi) \leq (|V| - 1)W$ (respectively, $\text{TP}(\pi) \geq -(|V| - 1)W$).*

Proof We prove the part for Min, the other case is similar. The proof proceeds by induction on the size of a partial play $\pi = v_1 \cdots v_k$ with $v_1 = v$. If $k \leq |V|$ then $\text{TP}(\pi) = \sum_{i=1}^{k-1} \omega(v_i, v_{i+1}) \leq (k-1)W \leq (|V|-1)W$. If $k \geq |V|+1$ then there exists $i < j$ such that $v_i = v_j$. Assume by contradiction that $\text{TP}(v_i \cdots v_j) > 0$. Then the play $\pi' = v_1 \cdots v_i \cdots v_j(v_{i+1} \cdots v_j)^\omega$ conforms to σ_{Min} and $\text{TP}(\pi') = +\infty$ which contradicts $\text{Val}(v, \sigma_{\text{Min}}) \neq +\infty$. Therefore $\text{TP}(v_i \cdots v_j) \leq 0$. We have $\text{TP}(\pi) = \text{TP}(v_1 \cdots v_i) + \text{TP}(v_i \cdots v_j) + \text{TP}(v_{j+1} \cdots v_k)$, and since $v_i = v_j$, $v_1 \cdots v_i v_{j+1} \cdots v_k$ is a finite play starting from v that conforms to σ_{Min} , and by induction hypothesis $\text{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) \leq (|V| - 1)W$. Then $\text{TP}(\pi) = \text{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) + \text{TP}(v_i \cdots v_j) \leq \text{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) \leq (|V| - 1)W$. $\square$

This permits to bound the finite values $\text{Val}(v)$ of vertices v of the game:

Corollary 24 *For all $v \in V$, $\text{Val}(v) \in [-(|V| - 1)W, (|V| - 1)W] \uplus \{-\infty, +\infty\}$.*

Proof From the result of Gimbert and Zielonka (2004), we know that total-payoff games are memorylessly determined, i.e. there exists two memoryless strategies $\sigma_{\text{Max}}, \sigma_{\text{Min}}$ such that for all v , $\text{Val}(v) = \text{Val}(v, \sigma_{\text{Max}}) = \text{Val}(v, \sigma_{\text{Min}})$. Assume that $\text{Val}(v) \notin \{-\infty, +\infty\}$. Then since $\text{Val}(v, \sigma_{\text{Min}}) = \text{Val}(v) \neq -\infty$, Lemma 23 shows that all finite play π that conforms to σ_{Min} verifies $\text{TP}(\pi) \geq -(|V| - 1)W$, therefore $\text{Val}(v) \geq -(|V| - 1)W$. One can similarly prove that $\text{TP}(v) \leq (|V| - 1)W$. $\square$

We now compare values in both games. A first lemma shows, in particular, that $\text{Val}_{\mathcal{G}^n}(v, n) \leq \text{Val}_{\mathcal{G}}(v)$, in case $\text{Val}_{\mathcal{G}}(v) \neq +\infty$.

Lemma 25 *For all $m \in \mathbb{Z}$, $v \in V$, and $n \geq 1$, if $\text{Val}_{\mathcal{G}}(v) \leq m$ then $\text{Val}_{\mathcal{G}^n}(v, n) \leq m$.*

Proof By hypothesis and using the memoryless determinacy of total-payoff games (Gimbert and Zielonka, 2004), there exists a memoryless strategy σ_{Min} for Min in $\mathcal{G}$ such that $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Min}}) \leq m$. Let σ_{Min}^m be the strategy in $\mathcal{G}^n$ defined, for all finite play π , vertex v' and $1 \leq j \leq n$, by

$$\begin{aligned} \sigma_{\text{Min}}^m(\pi(v', j)) &= (\mathbf{in}, \sigma_{\text{Min}}(\text{proj}(\pi)v'), j), \\ \sigma_{\text{Min}}^m(\pi(\mathbf{in}, v', j)) &= \begin{cases} (v', j) & \text{if } \text{TP}(\pi(\mathbf{in}, v', j)) \geq m + 1, \\ (\mathbf{ex}, v', j) & \text{if } \text{TP}(\pi(\mathbf{in}, v', j)) \leq m. \end{cases} \end{aligned}$$

Intuitively σ_{Min}^m simulates σ_{Min} , and asks to leave the copy when the current total-payoff is less than or equal to m . Notice that, by construction of σ_{Min}^m , $\text{proj}(\pi)$ conforms to σ_{Min} , if π conforms to σ_{Min}^m .

As a first step, if a play π starting in (v, n) and conforming to σ_{Min}^m encounters the target then its value is at most m . Indeed, it is of the form $\pi = \pi'(\text{in}, v', j)(\text{ex}, v', j)\mathbf{t}^\omega$, and since it conforms to σ_{Min}^m we have

$$\mathbf{MCR}(\pi) = \mathbf{TP}(\pi'(\text{in}, v', j)(\text{ex}, v', j)\mathbf{t}) = \mathbf{TP}(\pi'(\text{in}, v', j)) \leq m.$$

Then, assume, by contradiction, that there exists a play π starting in (v, n) and conforming to σ_{Min}^m , that does not encounter the target. Then, this means that Min does not ask $n+1$ times the ability to exit in π (since on the $(n+1)$ th time that we jump in an exterior vertex, Max is forced to go to the target). In particular, there exists $0 \leq j \leq n$ such that π is of the form $\pi'(v_1, j)(\text{in}, v_2, j)(v_2, j)(\text{in}, v_3, j) \cdots (v_k, j)(\text{in}, v_k, j) \cdots$. Since for all i , $\sigma_{\text{Min}}^m(\pi'(v_1, j)(\text{in}, v_2, j) \cdots (\text{in}, v_i, j)) = (v_i, j)$, we have that $\mathbf{TP}(\pi'(v_1, j)(\text{in}, v_2, j) \cdots (\text{in}, v_i, j)) \geq m+1$. Therefore, since every prefix of $\text{proj}(\pi)$ is the projection of a prefix of π , Lemma 22 shows that $\mathbf{TP}(\text{proj}(\pi)) \geq m+1 > m$, which raises a contradiction since $\text{proj}(\pi)$ conforms to σ_{Min} and $\text{Val}(v, \sigma_{\text{Min}}) \leq m$. Hence every play that conforms to σ_{Min}^m encounters the target, and, hence, has value at most m . This implies that $\text{Val}_{\mathcal{G}^n}(v, n) \leq m$. $\square$

We now turn to the other comparison between $\text{Val}_{\mathcal{G}^n}(v, n)$ and $\text{Val}_{\mathcal{G}}(v)$. Since $\text{Val}_{\mathcal{G}}(v)$ can be infinite in case the target is not reachable, we have to be more careful. In particular, we show that $\text{Val}_{\mathcal{G}^n}(v, n) \geq \min(\text{Val}_{\mathcal{G}}(v), (|V| - 1)W + 1)$ holds for large values of n . In the following, we let $K = |V|(2(|V| - 1)W + 1)$.

Lemma 26 *For all $m \leq (|V| - 1)W + 1$, $k \geq K$, and vertex v , if $\text{Val}_{\mathcal{G}}(v) \geq m$ then $\text{Val}_{\mathcal{G}^k}(v, k) \geq m$.*

Proof By hypothesis and using the memoryless determinacy proved by Gimbert and Zielonka (2004), there exists a memoryless strategy σ_{Max} for Max in $\mathcal{G}$ such that $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Max}}) \geq m$. Let σ_{Max}^m be the strategy in $\mathcal{G}^K$ defined, for all finite play π , vertex v' and $1 \leq j \leq n$, by:

$$\begin{aligned} \sigma_{\text{Max}}^m(\pi(v, j)) &= (\text{in}, \sigma_{\text{Max}}(\text{proj}(\pi)v), j), \\ \sigma_{\text{Max}}^m(\pi(\text{ex}, v, j)) &= \begin{cases} (v, j-1) & \text{if } \mathbf{TP}(\pi) \leq m-1 \text{ and } j > 1, \\ \mathbf{t} & \text{otherwise.} \end{cases} \end{aligned}$$

Intuitively σ_{Max}^m simulates σ_{Max} , and accepts to go to the target when the current total-payoff is greater than or equal to m .

By construction of σ_{Max}^m , if π conforms to σ_{Max}^m , then $\text{proj}(\pi)$ conforms to σ_{Max} . From the structure of the weighted graph, we know that for every play π of $\mathcal{G}^k$, there exists $1 \leq j \leq k$ such that π is of the form $\pi_k(\text{ex}, v_k, k)\pi_{k-1}(\text{ex}, v_{k-1}, k-1) \cdots \pi_j(\text{ex}, v_j, j)\pi'$ verifying that: there are no occurrences of exterior vertices in $\pi_k, \dots, \pi_j, \pi'$; for all $\ell \leq j$, all vertices in π_ℓ belong to the ℓ -th copy of $\mathcal{G}$; either $\pi' = \mathbf{t}^\omega$ or all vertices of π' belong to the $(j+1)$ th copy of $\mathcal{G}$ (in which case, $j < k$).

We now show that, in $\mathcal{G}^k$, $\mathbf{MCR}(\pi) \geq m$ for all play π starting in (v, k) and conforming to σ_{Max}^m . There are three cases to consider.

1. If π does not reach the target, then $\mathbf{MCR}(\pi) = +\infty \geq m$.
2. If $\pi = \pi_k(\text{ex}, v_k, k) \cdots \pi_j(\text{ex}, v_j, j)\mathbf{t}^\omega$ and $j > 1$ then,

$$\sigma_{\text{Max}}^m(\pi_k(\text{ex}, v_k, k) \cdots \pi_j(\text{ex}, v_j, j)) = \mathbf{t}.$$

Thus, using Lemma 22,

$$\begin{aligned}\mathbf{MCR}(\pi) &= \mathbf{TP}(\pi_k(\mathbf{ex}, v_k, k) \cdots \pi_j(\mathbf{ex}, v_j, j) \mathbf{t}) \\ &= \mathbf{TP}(\mathbf{proj}(\pi_k(\mathbf{ex}, v_k, k) \cdots \pi_j(\mathbf{ex}, v_j, j) \mathbf{t})) \\ &\geq \mathbf{Val}_{\mathcal{G}}(v, \sigma_{\mathbf{Max}}) \geq m.\end{aligned}$$

3. If $\pi = \pi_k(\mathbf{ex}, v_k, k) \cdots \pi_1(\mathbf{ex}, v_1, 1) \mathbf{t}^\omega$, assume by contradiction that

$$\mathbf{TP}(\pi_k(\mathbf{ex}, v_k, k) \cdots \pi_1) \leq m - 1.$$

Otherwise, we directly obtain $\mathbf{MCR}(\pi) \geq m$. Let v^* be a vertex that occurs at least $N = \lceil K/|V| \rceil = 2(|V| - 1)W + 1$ times in the sequence $v_1, \dots, v_k$: such a vertex exists, since otherwise $K \leq k \leq (N - 1)|V|$ which contradicts the fact that $(N - 1)|V| < K$. Let $j_1 > \dots > j_N$ be a sequence of indices such that $v_{j_i} = v^*$ for all i . We give a new decomposition of π :

$$\pi = \pi'_1(\mathbf{ex}, v_{j_1}, j_1) \cdots \pi'_N(\mathbf{ex}, v_{j_N}, j_N) \pi'_{N+1}.$$

Since π conforms to $\sigma_{\mathbf{Max}}^m$ and according to the assumption, we have that for all i ,

$$\mathbf{TP}(\pi'_1(\mathbf{ex}, v_{j_1}, j_1) \cdots \pi'_i) \leq m - 1.$$

We consider two cases.

- (a) If there exists π'_i such that $\mathbf{TP}(\pi'_i) \leq 0$ then, let $\mathbf{proj}(\pi'_i) = u_1 \cdots u_\ell$ with $u_1 = u_\ell = v^*$. Since π'_i conforms to $\sigma_{\mathbf{Max}}^m$, $\mathbf{proj}(\pi'_i)$ conforms to $\sigma_{\mathbf{Max}}$. Therefore the play

$$\tilde{\pi} = \mathbf{proj}(\pi'_1(\mathbf{ex}, v_{j_1}, j_1) \cdots \pi'_i(\mathbf{ex}, v_{j_i}, j_i))(u_1 \cdots u_{\ell-1})^\omega$$

conforms to $\sigma_{\mathbf{Max}}$. Furthermore, using again Lemma 22,

$$\mathbf{TP}(\tilde{\pi}) = \liminf_{n \rightarrow +\infty} \left(\mathbf{TP}(\pi'_1(\mathbf{ex}, v_{j_1}, i_1) \cdots \pi'_i(\mathbf{ex}, v_{j_i}, j_i)) + n \mathbf{TP}(u_1 \cdots u_\ell) \right)$$

and since $\mathbf{TP}(u_1 \cdots u_\ell) = \mathbf{TP}(\pi'_i) \leq 0$, we have

$$\mathbf{TP}(\tilde{\pi}) \leq \mathbf{TP}(\pi'_1(\mathbf{ex}, v_{j_1}, i_1) \cdots \pi'_i(\mathbf{ex}, v_{j_i}, j_i)) \leq m - 1.$$

Thus $\tilde{\pi}$ is a play starting from v that conforms to $\sigma_{\mathbf{Max}}$ but whose total-payoff is strictly less than m , which raises a contradiction.

- (b) If for all π'_i , $\mathbf{TP}(\pi'_i) \geq 1$ (notice that it is implied by $\mathbf{TP}(\pi'_i) > 0$). From Lemma 23, since $\mathbf{Val}_{\mathcal{G}}(v, \sigma_{\mathbf{Max}}) \geq m \neq -\infty$, we know that $\mathbf{TP}(\mathbf{proj}(\pi'_0)) \geq -(|V| - 1)W$. From Lemma 22, $\mathbf{TP}(\pi'_0) \geq -(|V| - 1)W$. Therefore

$$\begin{aligned}\mathbf{TP}(\pi'_1(\mathbf{ex}, v_{j_1}, i_1) \cdots \pi'_N) &\geq -(|V| - 1)W + N \\ &= (|V| - 1)W + 1 \geq m\end{aligned}$$

which contradicts the assumption that

$$\mathbf{TP}(\pi'_1(\mathbf{ex}, v_{j_1}, i_1) \cdots \pi'_N) < m.$$

We have shown that $\mathbf{MCR}(\pi) \geq m$ for all play π starting in (v, k) and conforming to $\sigma_{\mathbf{Max}}^m$, which implies $\mathbf{Val}_{\mathcal{G}^K}((v, k), \sigma_{\mathbf{Max}}^m) \geq m$. $\square$

Fig. 5 MCR game $\mathcal{G}_Y$ associated with the total-payoff game of Fig. 2

Using the two last lemmas, we can now prove Proposition 21 by relating precisely values in $\mathcal{G}$ and $\mathcal{G}^k$.

Proof (of Proposition 21) Let $v \in V$ be a vertex. We consider three cases:

- If $\text{Val}_{\mathcal{G}}(v) = -\infty$, then for all m , $\text{Val}_{\mathcal{G}}(v) \leq m$. By Lemma 25, $\text{Val}_{\mathcal{G}^K}(v, K) \leq m$. Therefore $\text{Val}_{\mathcal{G}^K}(v, K) = -\infty$.
- If $\text{Val}_{\mathcal{G}}(v) = m \in [-(|V| - 1)W, (|V| - 1)W]$. Then, $m \leq \text{Val}_{\mathcal{G}}(v) \leq m$. Thus, by Lemmas 25 and 26, $m \leq \text{Val}_{\mathcal{G}^K}(v, K) \leq m$. Therefore $\text{Val}_{\mathcal{G}^K}(v, K) = m$.
- If $\text{Val}_{\mathcal{G}}(v) = +\infty$, then $\text{Val}_{\mathcal{G}}(v) \geq (|V| - 1)W + 1$. By Lemma 26, $\text{Val}_{\mathcal{G}^K}(v, K) \geq (|V| - 1)W + 1$. $\square$

4.2 Value iteration algorithm for total-payoff games

By Proposition 21, an immediate way to obtain a value iteration algorithm for total-payoff games is to build game $\mathcal{G}^K$, run Algorithm 1 on it, and map the computed values back to $\mathcal{G}$. We take advantage of the structure of $\mathcal{G}^K$ to provide a better algorithm that avoids building $\mathcal{G}^K$. Intuitively, we first compute the values of the vertices in the *last copy of the game* (vertices of the form $(v, 1)$, $(\text{in}, v, 1)$ and $(\text{ex}, v, 1)$), then of those in the penultimate (vertices of the form $(v, 2)$, $(\text{in}, v, 2)$ and $(\text{ex}, v, 2)$), and so on.

We first sketch the intuitions that lead to the formalisation of these ideas. Let Z^j be a vector mapping each vertex v of $\mathcal{G}$ to the value $Z^j(v)$ of vertex (v, j) in $\mathcal{G}^K$. Then, we define an operator $\mathcal{H}$ such that $Z^{j+1} = \mathcal{H}(Z^j)$. Thus, assuming that we have computed the values of all vertices in the j th copy, $\mathcal{H}$ returns the value of the vertices in the $j+1$ st copy. In other words, the definition of $\mathcal{H}(Y)$ for some vector Y is to extract from $\mathcal{G}^K$ one copy of the game (that we call $\mathcal{G}_Y$), and make Y appear in the weights of some edges as illustrated in Fig. 5. This game, $\mathcal{G}_Y$, simulates a play in $\mathcal{G}$ in which Min can opt for ‘leaving the game’ at each round (by moving to the target), obtaining $\max(0, Y(v))$, if v is the current vertex. Then, we can define $\mathcal{H}(Y)(v)$ as the value of v in $\mathcal{G}_Y$. By construction, it is easy to see that $Z^{j+1} = \mathcal{H}(Z^j)$ holds for all $j \geq 1$, i.e. that $\mathcal{H}$ indeed corresponds to computing the values of the vertices in the $j+1$ st copy, given the values in the j th copy. Furthermore, letting $Z^0(v) = -\infty$ for all v , and $Z^1 = \mathcal{H}(Z^0)$, we will prove that: (i) $\mathcal{H}$ is monotonic, but may not be Scott-continuous; (ii) the sequence $(Z^j)_{j \geq 0}$ converges towards $\text{Val}_{\mathcal{G}}$. These arguments are the main ideas justifying Algorithm 2 to solve total-payoff games. Intuitively, the outer loop computes, in variable Y , a

Algorithm 2: A value iteration algorithm for total-payoff games

Input: Total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP} \rangle$, W largest weight in absolute value

```

1 foreach  $v \in V$  do  $Y(v) := -\infty$ 
2 repeat
3   foreach  $v \in V$  do  $Y_{pre}(v) := Y(v)$ ;  $Y(v) := \max(0, Y(v))$ ;  $X(v) := +\infty$ 
4   repeat
5      $X_{pre} := X$ 
6     foreach  $v \in V_{\text{Max}}$  do  $X(v) := \max_{v' \in E(v)} [\omega(v, v') + \min(X_{pre}(v'), Y(v'))]$ 
7     foreach  $v \in V_{\text{Min}}$  do  $X(v) := \min_{v' \in E(v)} [\omega(v, v') + \min(X_{pre}(v'), Y(v'))]$ 
8     foreach  $v \in V$  such that  $X(v) < -(|V| - 1)W$  do  $X(v) := -\infty$ 
9   until  $X = X_{pre}$ 
10   $Y := X$ 
11  foreach  $v \in V$  such that  $Y(v) > (|V| - 1)W$  do  $Y(v) := +\infty$ 
12 until  $Y = Y_{pre}$ 
13 return  $Y$ 

```

non-decreasing sequence of vectors whose limit is $\text{Val}_{\mathcal{G}}$, and that is stationary (this is not necessarily the case for the sequence $(Z^j)_{j \geq 0}$). Line 1 initialises Y to Z^0 . Each iteration of the outer loop amounts to running Algorithm 1 to compute $\mathcal{H}(Y_{pre})$ (lines 3 to 10), then detecting if some vertices have value $+\infty$, updating Y accordingly (line 11, following the second item of Proposition 21). We will show that, for all $j > 0$, if we let Y^j be the value of Y after the j -th iteration of the main loop, then $Z^j \preceq Y^j \preceq \text{Val}_{\mathcal{G}}$, which ensures the correctness of the algorithm.

Formal proof of the correctness of Algorithm 2 Let us now formalise the arguments sketched above to establish the correctness of Algorithm 2.

Theorem 27 *If a total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP} \rangle$ is given as input, Algorithm 2 outputs the vector $\text{Val}_{\mathcal{G}}$ of optimal values, after at most $K = |V|(2(|V| - 1)W + 1)$ iterations of the external loop. The complexity of the algorithm is $O(|V|^4|E|W^2)$.*

We first define formally the game $\mathcal{G}_Y$ described informally above. With the original total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP} \rangle$ and with every vector $Y \in \mathbb{Z}_{\infty}^V$, we associate the MCR game $\mathcal{G}_Y = \langle V', E_Y, \omega_Y, \{\mathbf{t}\}\text{-MCR} \rangle$ as follows. The sets of vertices are given by

$$V'_{\text{Max}} = V_{\text{Max}} \uplus \{\mathbf{t}\} \quad \text{and} \quad V'_{\text{Min}} = V_{\text{Min}} \uplus \{(\mathbf{in}, v) \mid v \in V\}.$$

As in game $\mathcal{G}^j$, vertices of the form $(\mathbf{in}, v)$ are called *interior vertices*. Edges are defined by

$$E_Y = \{(v, (\mathbf{in}, v')) \mid (v, v') \in E\} \uplus \{((\mathbf{in}, v), v) \mid v \in V\} \\ \uplus \{((\mathbf{in}, v), \mathbf{t}) \mid v \in V \wedge Y(v) \neq +\infty\} \uplus \{(\mathbf{t}, \mathbf{t})\}$$

while weights of edges are defined, for all $(v, v') \in E$, by

$$\omega_Y(v, (\mathbf{in}, v')) = \omega(v, v'), \quad \omega_Y((\mathbf{in}, v), \mathbf{t}) = \max(0, Y(v)), \\ \omega_Y((\mathbf{in}, v), v) = 0.$$

It is easy to see that lines 3 to 10 are a rewriting of Algorithm 1 in the special case of game $\mathcal{G}_Y$: in particular, neither the target vertex nor interior vertices are

explicit, but their behaviour is taken into account by the transformation performed in line 3 and the operators $\min$ used in the inner computation of lines 6 and 7. Hence, if we define $\mathcal{H}(Y)(v) = \text{Val}_{\mathcal{G}_Y}(v)$ for all $v \in V$, we can say that if inside the main loop, at line 3 the variable Y has value Y , then after line 10, it has value $\mathcal{H}(Y)$.

Notice that the game $\mathcal{G}_Y$ resembles a copy of $\mathcal{G}$ in the game $\mathcal{G}^j$ of the previous section. More, precisely, from the values $(\text{Val}_{\mathcal{G}^j}(v, j))_{v \in V}$ in the j th copy, we can deduce the values in the $(j + 1)$ th copy by an application of operator $\mathcal{H}$:

$$(\text{Val}_{\mathcal{G}^{j+1}}(v, j + 1))_{v \in V} = \mathcal{H}((\text{Val}_{\mathcal{G}^j}(v, j))_{v \in V}). \quad (5)$$

Although the 0th copy is not defined, we abuse the notation and set $\text{Val}_{\mathcal{G}^0}(v, 0) = -\infty$, which still conforms to the above equality. Furthermore, due to the structure of the game $\mathcal{G}^j$ notice that for all $j \leq j'$, $\text{Val}_{\mathcal{G}^j}(v, j) = \text{Val}_{\mathcal{G}^{j'}}(v, j)$.

Notice the absence of exterior vertices (ex, v', j) in game $\mathcal{G}_Y$, replaced by the computation of the maximum between 0 and $X(v')$ on the edge towards the target. Before proving the correctness of Algorithm 2, we prove several interesting properties of operator $\mathcal{H}$.

Proposition 28 *$\mathcal{H}$ is a monotonic operator.*

Proof For every vector $Y \in \mathbb{Z}_{\infty}^V$, let $\mathcal{F}_Y$ be the operator associated with the MCR game as defined in Section 3, i.e. for all $X \in \mathbb{Z}_{\infty}^{V'}$, and for all $v_1 \in V'$

$$\mathcal{F}_Y(X)(v_1) = \begin{cases} \max_{v_2 \in E_Y(v_1)} (\omega_Y(v_1, v_2) + X(v_2)) & \text{if } v_1 \in V'_{\text{Max}} \setminus \{\mathbf{t}\} \\ \min_{v_2 \in E_Y(v_1)} (\omega_Y(v_1, v_2) + X(v_2)) & \text{if } v_1 \in V'_{\text{Min}} \\ 0 & \text{if } v_1 = \mathbf{t}. \end{cases}$$

We know from Corollary 11 that $\text{Val}_{\mathcal{G}_Y}$ is the greatest fixed point of $\mathcal{F}_Y$. Consider now two vectors $Y, Y' \in \mathbb{Z}_{\infty}^V$ such that $Y \preceq Y'$.

First, notice that for all $X \in \mathbb{Z}_{\infty}^{V'}$:

$$\mathcal{F}_Y(X) \preceq \mathcal{F}_{Y'}(X). \quad (6)$$

Indeed, to get the result it suffices to notice that for all $v_1, v_2 \in V'$, $\omega_Y(v_1, v_2) \leq \omega_{Y'}(v_1, v_2)$ (the only differences are on the edges $((\text{in}, v), \mathbf{t})$).

Consider then the vector X_0 defined by $X_0(v_1) = +\infty$ for all $v_1 \in V'$. Given $Y \preceq Y'$, from (6), we have that $\mathcal{F}_Y(X_0) \preceq \mathcal{F}_{Y'}(X_0)$, then a simple induction shows that for all i , $\mathcal{F}_Y^i(X_0) \preceq \mathcal{F}_{Y'}^i(X_0)$. Thus, since $\text{Val}_{\mathcal{G}_Y}$ (respectively, $\text{Val}_{\mathcal{G}_{Y'}}$) is the greatest fixed point of $\mathcal{F}_Y$ (respectively, $\mathcal{F}_{Y'}$), we have $\text{Val}_{\mathcal{G}_Y} \preceq \text{Val}_{\mathcal{G}_{Y'}}$. As a consequence $\mathcal{H}(Y) \preceq \mathcal{H}(Y')$. $\square$

Notice that $\mathcal{H}$ may not be Scott-continuous, as shown in the following example.

Example 29 Recall that, in our setting, a Scott-continuous operator is a mapping $F : \mathbb{Z}_{\infty}^V \rightarrow \mathbb{Z}_{\infty}^V$ such that for every sequence of vectors $(x_i)_{i \geq 0}$ having a limit x_{ω} , the sequence $(F(x_i))_{i \geq 0}$ has a limit equal to $F(x_{\omega})$.

We present a total-payoff game whose associated operator $\mathcal{H}$ is not continuous. Let $\mathcal{G}$ be the total-payoff game containing one vertex v of Min and a self loop of weight -1 (as depicted in Fig. 6). For all $Y \in \mathbb{Z}$, in the MCR game $\mathcal{G}_Y$, v

Fig. 6 An example where the operator $\mathcal{H}$ is not Scott-continuous.

has value $-\infty$, indeed one can take the loop an arbitrary number of times before reaching the target, ensuring a value arbitrary low. Therefore, if we take an increasing sequence $(Y_i)_{i \geq 0}$ of integers, $\mathcal{H}(Y_i)(v) = -\infty$ for all i , thus the limit of the sequence $(\mathcal{H}(Y_i))_{i \geq 0}$ is $-\infty$. However, the limit of the sequence $(Y_i)_{i \geq 0}$ is $+\infty$ and $\mathcal{H}(+\infty)(v) = +\infty$, since the target is not reachable anymore (in case the weight of an edge would be $+\infty$, it is removed in the definition of E_Y). Thus, $\mathcal{H}$ is not Scott-continuous. $\square$

In particular, we may not use the Kleene sequence, as we have done for MCR games, to conclude to the correctness of our algorithm. Anyhow, we will show that the sequence $(Y^j)_{j \geq 0}$ indeed converges towards the vector of values of the total-payoff game. We first show that this vector is a pre-fixed point of $\mathcal{H}$ starting with a technical lemma that is useful in the subsequent proof.

Lemma 30 *Let σ_{Min} be a strategy for Min in $\mathcal{G}$, and $\pi = v_1 \cdots v_i$ a finite play that conforms to σ_{Min} . Then:*

$$\mathbf{TP}(v_1 \cdots v_i) + \mathbf{Val}_{\mathcal{G}}(v_i) \leq \mathbf{Val}_{\mathcal{G}}(v_1, \sigma_{\text{Min}}).$$

Proof Let σ_{Max} be an optimal strategy for Max and $v_i v_{i+1} v_{i+2} \cdots$ be the play $\text{Play}(v_i, \sigma_{\text{Max}}, \sigma_{\text{Min}})$. Since σ_{Max} is optimal, $\mathbf{TP}(v_i v_{i+1} v_{i+2} \cdots) \geq \mathbf{Val}_{\mathcal{G}}(v_i)$. Furthermore notice that $v_1 \cdots v_i v_{i+1} \cdots$ conforms to the strategy σ_{Min} , therefore we have $\mathbf{TP}(v_1 v_1 \cdots v_i v_{i+1} \cdots) \leq \mathbf{Val}_{\mathcal{G}}(v_1, \sigma_{\text{Min}})$. Thus:

$$\begin{aligned} \mathbf{TP}(v_1 \cdots v_i) + \mathbf{Val}_{\mathcal{G}}(v_i) &\leq \mathbf{TP}(v_1 \cdots v_i) + \mathbf{TP}(v_i v_{i+1} \cdots) \\ &\leq \mathbf{TP}(v_1 v_1 \cdots v_i v_{i+1} \cdots) \leq \mathbf{Val}_{\mathcal{G}}(v_1, \sigma_{\text{Min}}). \end{aligned} \quad \square$$

Lemma 31 *$\mathbf{Val}_{\mathcal{G}}$ is a pre-fixed point of $\mathcal{H}$, i.e. $\mathcal{H}(\mathbf{Val}_{\mathcal{G}}) \preceq \mathbf{Val}_{\mathcal{G}}$.*

Proof To ease the notations, we denote $\mathbf{Val}_{\mathcal{G}}$ by Y^* in this proof. To prove this lemma, we just have to show that for all $v_1 \in V$, the value of v_1 in the MCR game $\mathcal{G}_{Y^*}$ is at most its value in the original total-payoff game $\mathcal{G}$, i.e.

$$\mathcal{H}(Y^*)(v_1) = \mathbf{Val}_{\mathcal{G}_{Y^*}}(v_1) \leq Y^*(v_1).$$

Let σ_{Min} be a memoryless strategy in $\mathcal{G}$ such that $\mathbf{Val}_{\mathcal{G}}(v_1, \sigma_{\text{Min}}) \leq m$ for some $m \in \mathbb{Z} \uplus \{+\infty\}$. And let σ_{Min}^m be a strategy in $\mathcal{G}_X$ defined for all finite play $\pi v'$ with $v' \in V_{\text{Min}}$ by $\sigma_{\text{Min}}^m(\pi v') = (\text{in}, \sigma_{\text{Min}}(v'))$ and for all finite play $\pi(\text{in}, v')$,

$$\sigma_{\text{Min}}^m(\pi(\text{in}, v')) = \begin{cases} \mathbf{t} & \text{if } \mathbf{TP}(\pi(\text{in}, v')\mathbf{t}) \leq m \\ v' & \text{otherwise.} \end{cases}$$

Notice that, by construction, all plays starting in v_1 , conforming to σ_{Min}^m and that reach the target have value at most m . Assume by contradiction that there exists a play $v_1(\text{in}, v_2)v_2(\text{in}, v_3) \cdots \in \text{Play}(v_1, \sigma_{\text{Min}}^m)$ that never reaches $\mathbf{t}$. In particular, for all i , $\mathbf{TP}(v_1(\text{in}, v_2) \cdots (\text{in}, v_i)\mathbf{t}) > m$.

Again by construction, $v_1v_2 \cdots$ is a play in $\mathcal{G}$ that conforms to σ_{Min} and for all i , $\mathbf{TP}(v_1(\text{in}, v_2) \cdots (\text{in}, v_i)) = \mathbf{TP}(v_1 \cdots v_i)$.

- If there exists $i \geq 2$ such that $Y^*(v_i) = \text{Val}_{\mathcal{G}}(v_i) \geq 0$, then

$$\begin{aligned} \mathbf{TP}(v_1(\text{in}, v_2) \cdots (\text{in}, v_i)\mathbf{t}) &= \mathbf{TP}(v_1(\text{in}, v_2) \cdots (\text{in}, v_i)) + \omega_X((\text{in}, v_i), \mathbf{t}) \\ &= \mathbf{TP}(v_1 \cdots v_i) + \max(0, Y^*(v_i)) \\ &= \mathbf{TP}(v_1 \cdots v_i) + Y^*(v_i) \\ &\leq \text{Val}(v_1, \sigma_{\text{Min}}) \quad (\text{from Lemma 30}) \\ &\leq m \end{aligned}$$

which raises a contradiction.

- If for all $i \geq 2$, $Y^*(v_i) = \text{Val}_{\mathcal{G}}(v_i) < 0$ then for all $i \geq 2$,

$$\mathbf{TP}(v_1(\text{in}, v_2) \cdots (\text{in}, v_i)\mathbf{t}) = \mathbf{TP}(v_1(\text{in}, v_2) \cdots (\text{in}, v_i)) = \mathbf{TP}(v_1 \cdots v_i) > m.$$

Thus $\mathbf{TP}(v_1v_2 \cdots) > m$, which contradicts the fact that $v_1v_2 \cdots$ conforms to σ_{Min} and $\text{Val}_{\mathcal{G}}(v_1, \sigma_{\text{Min}}) \leq m$.

Thus $\text{Val}_{\mathcal{G}_{Y^*}}(v_1, \sigma_{\text{Min}}) \leq m$. As a consequence, $\text{Val}_{\mathcal{G}_{Y^*}}(v_1) \leq \text{Val}_{\mathcal{G}}(v_1)$. $\square$

Remark 32 Even if it is not necessary for the proof of Theorem 27, we can show that $\text{Val}_{\mathcal{G}}$ is the least pre-fixed point of $\mathcal{H}$. Notice that, by monotonicity of $\mathcal{H}$, this directly implies that $\text{Val}_{\mathcal{G}}$ is the least fixed point of $\mathcal{H}$. The proof that $\text{Val}_{\mathcal{G}}$ is the least pre-fixed point of $\mathcal{H}$ amounts to better understand the convergence of the sequence $(\text{Val}_{\mathcal{G}^j}(v, j))_{j \geq 0}$ (remember that we set $\text{Val}_{\mathcal{G}^0}(v, 0) = -\infty$ for all v). Indeed, Proposition 21 already shows that for vertices v such that $\text{Val}_{\mathcal{G}}(v) < +\infty$, $(\text{Val}_{\mathcal{G}^j}(v, j))_{j \geq 1}$ converges towards $\text{Val}_{\mathcal{G}}(v)$. It is also the case for vertices v such that $\text{Val}_{\mathcal{G}}(v) = +\infty$ as we show in Lemma 33 below. Then, consider a pre-fixed point Y of $\mathcal{H}$, i.e. $\mathcal{H}(Y) \preceq Y$. Since $\text{Val}_{\mathcal{G}^0}(v, 0) = -\infty \leq Y(v)$ and $\mathcal{H}$ is monotonous, we prove by immediate induction that $\text{Val}_{\mathcal{G}^j}(v, j) \leq Y(v)$ for all v and $j \geq 0$: indeed, if $\text{Val}_{\mathcal{G}^j}(v, j) \leq Y(v)$ for all v , we have

$$(\text{Val}_{\mathcal{G}^{j+1}}(v, j+1))_{v \in V} = \mathcal{H}((\text{Val}_{\mathcal{G}^j}(v, j))_{v \in V}) \preceq \mathcal{H}(Y) \preceq Y.$$

This implies that $\text{Val}_{\mathcal{G}} \preceq Y$, showing that $\text{Val}_{\mathcal{G}}$ is indeed the least pre-fixed point of $\mathcal{H}$, and hence the least fixed point of $\mathcal{H}$, by the above reasoning.

Before continuing the proof of Theorem 27, we show the result used in the previous remark.

Lemma 33 *Let $v \in V$ such that $\text{Val}_{\mathcal{G}}(v) = +\infty$, and σ_{Max} a memoryless strategy for Max in $\mathcal{G}$ such that $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Max}}) = +\infty$. Then the following holds:*

- (i) *For every finite play $v_1 \cdots v_k$ conforming to σ_{Max} starting in $v_1 = v$, if there exists $i < j$ such that $v_i = v_j$ then $\mathbf{TP}(v_i \cdots v_j) \geq 1$.*
- (ii) *For every $m \in \mathbb{N}$, $k \geq m|V| + 1$ and $v_1 \cdots v_k$ a finite play conforming to σ_{Max} and starting in $v_1 = v$, $\mathbf{TP}(v_1 \cdots v_k) \geq m - (|V| - 1)W$.*

- (iii) For all $m \in \mathbb{N}$ and $k \geq (m + (|V| - 1)W)|V| + 1$, $\text{Val}_{\mathcal{G}^k}(v, k) \geq m$.
 (iv) $\lim_{j \rightarrow \infty} \text{Val}_{\mathcal{G}^j}(v, j) = +\infty$.

Proof We prove (i) by contradiction. Therefore, assume that $\text{TP}(v_i \cdots v_j) \leq 0$, then $\pi = v_1 \cdots v_{i-1}(v_i \cdots v_{j-1})^\omega$ conforms to σ_{Max} and $\text{TP}(\pi) \leq \text{TP}(v_1 \cdots v_{i-1}) < +\infty$, which contradicts the fact that $\text{Val}_{\mathcal{G}}(v, \sigma_{\text{Max}}) = +\infty$.

We prove (ii) by induction on m . The base case is straightforward. For the inductive case, let $m > 0$, $k \geq m|V| + 1$ and $v_1 \cdots v_k$ a finite play conforming to σ_{Max} and starting in $v_1 = v$. Since $k \geq m|V| + 1 \geq |V| + 1$ there exists $i < j \leq |V| + 1$ such that $v_i = v_j$. Thus:

$$\begin{aligned} \text{TP}(v_1 \cdots v_k) &= \text{TP}(v_1 \cdots v_i) + \text{TP}(v_i \cdots v_j) + \text{TP}(v_j \cdots v_k) \\ &= \text{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) + \text{TP}(v_i \cdots v_j) \\ &\geq \text{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) + 1. \end{aligned} \quad \text{from (i)}$$

By induction hypothesis, as $v_1 \cdots v_i v_{j+1} \cdots v_k$ conforms to σ_{Max} and has length at least $(m - 1)|V| + 1$ (because $i < j \leq |V| + 1$, implying that $j - i \leq |V|$), we have $\text{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) \geq m - 1 - (|V| - 1)W$, thus $\text{TP}(v_1 \cdots v_k) \geq m - (|V| - 1)W$.

To prove (iii), let σ'_{Max} be a strategy of Max in $\mathcal{G}^k$ defined by $\sigma'_{\text{Max}}(v, j) = (\text{in}, \sigma_{\text{Max}}(v, j))$ and $\sigma'_{\text{Max}}(\text{ex}, v, j) = (v, j - 1)$ for all $v \in V$ and $j \geq k$. Let π be a play starting in (v, k) and conforming to σ'_{Max} . If π does not reach $\mathfrak{t}$, then $\text{MCR}(\pi) = +\infty \geq m$. If π reaches the target then $\text{proj}(\pi)$ is of the form $v_1 \cdots v_\ell \mathfrak{t}^\omega$, with $\text{MCR}(\pi) = \text{TP}(v_1 \cdots v_\ell)$. It is clear by construction of σ'_{Max} that $v_1 \cdots v_\ell$ is a finite play of $\mathcal{G}$ that conforms to σ_{Max} . Furthermore, $\ell \geq k \geq (m + (|V| - 1)W)|V| + 1$ thus, from (ii), we have that $\text{TP}(v_1 \cdots v_\ell) \geq m$. This implies $\text{MCR}(\pi) \geq m$. Hence, every play in $\mathcal{G}^k$ conforming to σ'_{Max} and starting in (v, k) has a value at least m , which means that $\text{Val}_{\mathcal{G}^k}(v, k) \geq m$.

Item (iv) is then a direct consequence of (iii). $\square$

We are now ready to state and prove the inductive invariant allowing us to show the correctness of Algorithm 2.

Lemma 34 *Before the j -th iteration of the external loop of Algorithm 2, we have $\text{Val}_{\mathcal{G}^j}(v, j) \leq Y^j(v) \leq \text{Val}_{\mathcal{G}}(v)$ for all vertices $v \in V$.*

Proof For $j = 0$, we have $Y^0(v) = -\infty = \text{Val}_{\mathcal{G}^0}(v, 0)$ for all vertex $v \in V$. Suppose then that the invariant holds for $j \geq 0$. We know (see (5)) that $\text{Val}_{\mathcal{G}^{j+1}}(v, j + 1) = \mathcal{H}((\text{Val}_{\mathcal{G}^j}(v', j))_{v' \in V})$. Moreover, after the assignment of line 10, by definition of $\mathcal{H}$, variable Y contains $\mathcal{H}(Y^j)$. The operation performed on line 11 only increases the values of vector Y , so that at the end of the j th iteration, we have $\mathcal{H}(Y^j) \preceq Y^{j+1}$. Since $\mathcal{H}$ is monotonous, and by the invariant at step j , we obtain

$$\text{Val}_{\mathcal{G}^{j+1}}(v, j + 1) = \mathcal{H}((\text{Val}_{\mathcal{G}^j}(v', j))_{v' \in V}) \leq \mathcal{H}(Y^j) \leq Y^{j+1}.$$

Moreover, using again the monotony of $\mathcal{H}$ and Lemma 31, we have

$$\mathcal{H}(Y^j) \preceq \mathcal{H}(\text{Val}_{\mathcal{G}}) \preceq \text{Val}_{\mathcal{G}}.$$

A closer look at line 11 shows that $\mathcal{H}(Y^j)$ and Y^{j+1} coincide over vertices v such that $\mathcal{H}(Y^j)(v) \leq (|V| - 1)W$, and otherwise $Y^{j+1}(v) = +\infty$. Hence, if $\mathcal{H}(Y^j)(v) \leq (|V| - 1)W$, we directly obtain $Y^{j+1}(v) = \mathcal{H}(Y^j)(v) \leq \text{Val}_{\mathcal{G}}(v)$. Otherwise, we know

Fig. 7 An example total-payoff game where each execution of the inner loop requires two iterations while the outer loop takes W iterations to stabilise.

that $\text{Val}_{\mathcal{G}}(v) > (|V| - 1)W$. By Corollary 24, we know that $\text{Val}_{\mathcal{G}}(v) = +\infty$, so that $Y^{j+1}(v) = +\infty = \text{Val}_{\mathcal{G}}(v)$. In the overall, we have proved

$$\text{Val}_{\mathcal{G}^{j+1}}(v, j+1) \leq Y^{j+1}(v) \leq \text{Val}_{\mathcal{G}}(v) \quad \square$$

We are now able to prove the correctness and termination of the algorithm.

Proof (of Theorem 27) For $j = K$ (remember that K was defined in the previous section), the invariant of Lemma 34 becomes

$$\text{Val}_{\mathcal{G}^K}(v, K) \leq Y^K(v) \leq \text{Val}_{\mathcal{G}}(v)$$

for all vertices $v \in V$. Notice that the iteration may have stopped before iteration K , in which case the sequence $(Y^j)_{j \geq 0}$ may be considered as stationary. In case $\text{Val}_{\mathcal{G}}(v) \neq +\infty$, Proposition 21 proves that $\text{Val}_{\mathcal{G}^K}(v, K) = \text{Val}_{\mathcal{G}}(v)$, so that we have $Y^K(v) = \text{Val}_{\mathcal{G}}(v)$. In case $\text{Val}_{\mathcal{G}}(v) = +\infty$, Proposition 21 shows that $\text{Val}_{\mathcal{G}^K}(v, K) > (|V| - 1)W$: by the operation performed at line 11, we obtain that $Y^K(v) = +\infty = \text{Val}_{\mathcal{G}}(v)$.

Hence, $K = |V|(2(|V| - 1)W + 1)$ is an upper bound on the number of iterations before convergence of Algorithm 2, and moreover, at the convergence, the algorithm outputs the vector of optimal values of the total-payoff game. $\square$

Convergence speed of the algorithm Observe that the number of iterations in each internal loop is controlled by Theorem 3. On the example of Fig. 2, only 2 external iterations are necessary, but the number of iterations of each internal loop would be $2W$. By contrast, for the total-payoff game depicted in Fig. 7, each internal loop requires 2 iterations to converge, but the external loop takes W iterations to stabilise. A combination of both examples would experience a pseudo-polynomial number of iterations to converge in both the internal and external loops, matching the W^2 term of the above complexity: this gives rise to the parametric example of Fig. 8.

4.3 Optimal strategies

In Section 3, we have shown, for all MCR games, the existence of a fake-optimal NC-strategy permitting to reconstruct an optimal finite-memory strategy for Min (if every vertex has value different from $-\infty$, or a strategy ensuring every possible threshold for vertices with value $-\infty$). Given a total-payoff game $\mathcal{G}$, if we apply this construction to the game $\mathcal{G}_{\text{Val}_{\mathcal{G}}}$, we obtain an NC-strategy σ_{Min}^* . Consider the strategy $\overline{\sigma}_{\text{Min}}$, obtained by projecting σ_{Min}^* on V as follows: for all finite plays π and vertices $v \in V_{\text{Min}}$, we let $\overline{\sigma}_{\text{Min}}(\pi v) = v'$ if $\sigma_{\text{Min}}^*(v) = (\text{in}, v')$. We show thereafter that

Fig. 8 Parametric weighted graph

$\overline{\sigma_{\text{Min}}}$ is optimal for Min in $\mathcal{G}$. Notice that σ_{Min}^* , and hence $\overline{\sigma_{\text{Min}}}$, can be computed during the last iteration of the value iteration algorithm, as explained in the case of MCR games in Section 3.4. A similar construction can be done to compute an optimal strategy for Max.

Theorem 35 *The memoryless strategy $\overline{\sigma_{\text{Min}}}$ is optimal in $\mathcal{G}$.*

Proof We start by showing that for all v ,

$$(1) \text{ if } \sigma_{\text{Min}}^*(\text{in}, v) = \mathbf{t}, \text{ then } \text{Val}_{\mathcal{G}}(v) = \omega((\text{in}, v), \mathbf{t}) \geq 0.$$

By definition of σ_{Min}^* , $\mathbf{t} = \text{argmin}_{v' \in E(\text{in}, v)} (\omega((\text{in}, v), v') + \text{Val}_{\mathcal{G}_{\text{Val}_{\mathcal{G}}}}(v'))$. In particular, for $v' = v$, as $\omega((\text{in}, v), v) = 0$, we have $\text{Val}_{\mathcal{G}_{\text{Val}_{\mathcal{G}}}}(v) \geq \omega((\text{in}, v), \mathbf{t}) + \text{Val}_{\mathcal{G}_{\text{Val}_{\mathcal{G}}}}(\mathbf{t})$, i.e. $\text{Val}_{\mathcal{G}}(v) \geq \omega((\text{in}, v), \mathbf{t})$. Reciprocally, we know that $\omega((\text{in}, v), \mathbf{t}) = \max(\text{Val}_{\mathcal{G}}(v), 0) \geq \text{Val}_{\mathcal{G}}(v)$, so that, we finally have $\text{Val}_{\mathcal{G}}(v) = \omega((\text{in}, v), \mathbf{t})$.

We then show that

- (2) for all plays $\pi = v_0 v_1 \dots$ in $\mathcal{G}$ that conforms with $\overline{\sigma_{\text{Min}}}$ and such that $\text{Val}_{\mathcal{G}}(v_0) < +\infty$, either $\text{TP}(\pi) = -\infty$ or there exists $i_\pi > 0$ such that $\sigma_{\text{Min}}^*(\text{in}, v_{i_\pi}) = \mathbf{t}$ and $\text{TP}(v_0 \dots v_{i_\pi}) \leq \text{Val}_{\mathcal{G}}(v_0) - \text{Val}_{\mathcal{G}}(v_{i_\pi})$.

Assume first that during π , Min never asks to go to the target, i.e. for all $i > 0$, $\sigma_{\text{Min}}^*(\text{in}, v_i) = v_i$. Then the play $v_0(\text{in}, v_1)v_1(\text{in}, v_2) \dots$ is a play of $\mathcal{G}_{\text{Val}_{\mathcal{G}}}$ that conforms with σ_{Min}^* . As there are only finitely many vertices in $\mathcal{G}_{\text{Val}_{\mathcal{G}}}$, there must

exist a vertex v that appears infinitely often in this play. As σ_{Min}^* is an NC-strategy, the accumulated cost of the chunk of the play between two occurrences of v has weight at most -1 , thus the total payoff of $v_0(\text{in}, v_1)v_1(\text{in}, v_2)\dots$ is $-\infty$. As the total-payoff of this play is equal to the total payoff of π , we have that $\text{TP}(\pi) = -\infty$.

Otherwise, Min asks at some point to go to the target: let i_π be the first index such that $\sigma_{\text{Min}}^*(\text{in}, v_{i_\pi}) = \mathbf{t}$. As the strategy σ_{Min}^* is a fake-optimal NC-strategy, we know that the accumulated cost of π until the target verifies $\text{TP}(\pi) = \text{MCR}(\pi) \leq \text{Val}_{\mathcal{G}_{\text{Val}_{\mathcal{G}}}}(v_0)$. As $\text{Val}_{\mathcal{G}}$ is a fixed point of $\mathcal{H}$ (see Remark 32), we have that for all v , $\text{Val}_{\mathcal{G}_{\text{Val}_{\mathcal{G}}}}(v) = \mathcal{H}(\text{Val}_{\mathcal{G}}(v)) = \text{Val}_{\mathcal{G}}(v)$, thus $\text{TP}(\pi) \leq \text{Val}_{\mathcal{G}}(v_0)$. We have $\text{TP}(v_0 \dots v_{i_\pi}) = \text{TP}(v_0(\text{in}, v_1)v_1 \dots v_{i_\pi-1}(\text{in}, v_{i_\pi})) = \text{TP}(\pi) - \omega((\text{in}, v_{i_\pi}), \mathbf{t})$. From (1), we have $\omega((\text{in}, v_{i_\pi}), \mathbf{t}) = \text{Val}_{\mathcal{G}}(v_{i_\pi})$, thus $\text{TP}(v_0 \dots v_{i_\pi}) \leq \text{Val}_{\mathcal{G}}(v_0) - \text{Val}_{\mathcal{G}}(v_{i_\pi})$, which proves (2).

Now let us prove the theorem. Let v be a vertex in $\mathcal{G}$. If $\text{Val}_{\mathcal{G}}(v) = +\infty$ then trivially $\text{Val}_{\mathcal{G}}(v, \overline{\sigma}_{\text{Min}}) \leq +\infty = \text{Val}_{\mathcal{G}}(v)$. Otherwise let $\pi = v_0v_1\dots$ be a play in $\mathcal{G}$ that conforms with $\overline{\sigma}_{\text{Min}}$ such that $v = v_0$. If $\text{TP}(\pi) = -\infty$ then $\text{TP}(\pi) \leq \text{Val}_{\mathcal{G}}(v)$. Otherwise we construct inductively an increasing sequence of indices $i_0, i_1, \dots$ such that $\text{TP}(v_{i_j} \dots v_{i_{j+1}}) \leq \text{Val}_{\mathcal{G}}(v_{i_j}) - \text{Val}_{\mathcal{G}}(v_{i_{j+1}})$ (for all j) as follows. First, we let $i_0 = 0$. Then, for all j , let $\pi' = v_{i_j}v_{i_j+1}\dots$ be the current suffix of π : since $\text{TP}(\pi) \neq -\infty$, we have $\text{TP}(\pi') \neq -\infty$, thus (2) shows that by letting $i_{j+1} = i_j + i_{\pi'}$, we obtain $\text{TP}(v_{i_j} \dots v_{i_{j+1}}) \leq \text{Val}_{\mathcal{G}}(v_{i_j}) - \text{Val}_{\mathcal{G}}(v_{i_{j+1}})$, and we know that $\sigma_{\text{Min}}^*(\text{in}, v_{i_{j+1}}) = \mathbf{t}$.

We can then show, by a direct induction, that for all $j \geq 1$, $\text{TP}(v_0 \dots v_{i_j}) \leq \text{Val}_{\mathcal{G}}(v) - \text{Val}_{\mathcal{G}}(v_{i_j})$. From (1), we know that $\text{Val}_{\mathcal{G}}(v_{i_j}) \geq 0$, thus $\text{TP}(v_0 \dots v_{i_j}) \leq \text{Val}_{\mathcal{G}}(v)$. Hence

$$\text{TP}(\pi) = \liminf_{n \rightarrow \infty} \text{TP}(v_0 \dots v_n) \leq \liminf_{j \rightarrow \infty} \text{TP}(v_0 \dots v_{i_j}) \leq \text{Val}_{\mathcal{G}}(v). \quad \square$$

5 Implementation and heuristics

In this section, we report on a prototype implementation of our algorithms.⁷ For convenience reasons, we have implemented them as an add-on to PRISM-games (Chen et al, 2013), although we could have chosen to extend another model-checker as we do not rely on the probabilistic features of PRISM models (i.e. we use the PRISM syntax of *stochastic multi-player games*, allowing arbitrary rewards, and forbidding probability distributions different of Dirac ones). We then use rPATL specifications of the form $\langle\langle C \rangle\rangle R^{\min / \max=?} [\mathbf{F}^\infty \varphi]$ and $\langle\langle C \rangle\rangle R^{\min / \max=?} [\mathbf{F}^c \perp]$ to model respectively MCR games and total-payoff games, where C represents a coalition of players that want to minimise/maximise the payoff, and φ is another rPATL formula describing the target set of vertices (for total-payoff games, such a formula is not necessary). We have tested our implementation on toy examples. On the parametric one Fig. 8, results obtained by applying our algorithm for total-payoff games are summarised in Table 1, where for each pair (W, n) , we give the time t in seconds, the number k_e of iterations in the external loop, and the total number k_i of iterations in the internal loop.

⁷ Source and binary files, as well as some examples, can be downloaded from <http://www.ulb.ac.be/di/verif/monmege/tool/TP-MCR/>.

Table 1 Results of value iteration on a parametric example

W	n	without heuristics			with heuristics		
		t	k_e	k_i	t	k_e	k_i
50	100	0.52s	151	12,603	0.01s	402	1,404
50	500	9.83s	551	53,003	0.42s	2,002	7,004
200	100	2.96s	301	80,103	0.02s	402	1,404
200	500	45.64s	701	240,503	0.47s	2,002	7,004
500	1,000	536s	1,501	1,251,003	2.37s	4,002	14,004

Notice that due to the very little memory consumption of the algorithm, there is no risk of running out of memory. However, the execution time can become very large. For instance, in case $W = 500$ and $n = 1000$, the execution time becomes 536s whereas the total number of iterations in the internal loop is greater than a million.

5.1 Acceleration techniques

We close this section by sketching two techniques that can be used to speed up the computation of the fixed point in Algorithms 1 and 2. We fix a weighted graph $\langle V, E, \omega \rangle$. Both accelerations rely on a topological order of the strongly connected components (SCC for short) of the graph, given as a function $c: V \rightarrow \mathbb{N}$, mapping each vertex to its *component*, verifying that (i) $c(V) = \{0, \dots, p\}$ for some $p \geq 0$, (ii) $c^{-1}(q)$ is a maximal SCC for all q , (iii) and $c(v) \geq c(v')$ for all $(v, v') \in E$.⁸ In case of an MRC game with $\mathbf{t}$ the unique target, $c^{-1}(0) = \{\mathbf{t}\}$. Intuitively, c induces a directed acyclic graph whose vertices are the sets $c^{-1}(q)$ for all $q \in c(V)$, and with an edge (S_1, S_2) if and only if there are $v_1 \in S_1, v_2 \in S_2$ such that $(v_1, v_2) \in E$.

The *first acceleration heuristic* is a divide-and-conquer technique that consists in applying Algorithm 1 (or the inner loop of Algorithm 2) iteratively on each $c^{-1}(q)$ for $q = 0, 1, 2, \dots, p$, using at each step the information computed during steps $j < q$ (since the value of a vertex v depends only on the values of the vertices v' such that $c(v') \leq c(v)$).

The *second acceleration heuristic* consists in studying more precisely each component $c^{-1}(q)$. Having already computed the optimal values $\text{Val}(v)$ of vertices $v \in c^{-1}(\{0, \dots, q-1\})$, we ask an oracle to precompute a finite set $S_v \subseteq \mathbb{Z}_\infty$ of possible optimal values for each vertex $v \in c^{-1}(q)$. For MCR games and the inner iteration of the algorithm for total-payoff games, one way to construct such a set S_v is to consider that possible optimal values are the one of non-looping paths inside the component exiting it, since, in MCR games, there exist optimal strategies for both players whose outcome is a non-looping path (see Section 3).

Algorithm 3 and 4 are enhanced versions of Algorithm 1 and 2 respectively, that apply these acceleration heuristics.

Finally, we note that we can identify classes of weighted graphs for which there exists an oracle that runs in polynomial time and returns, for all vertices v , a set S_v of polynomial size. On such classes, Algorithms 1 and 2, enhanced with our two acceleration techniques, *run in polynomial time*. For instance, for all fixed positive integers L , the class of weighted graphs where every component $c^{-1}(q)$ uses at

⁸ Such a mapping is computable in linear time, e.g., by Tarjan's algorithm (Tarjan, 1972).

Algorithm 3: Accelerated value iteration algorithm for min-cost reachability games

Input: MCR game $\langle V, E, \omega, \{\mathbf{t}\}\text{-MCR} \rangle$, SCC-decomposition $c: V \rightarrow \{0, 1, \dots, p\}$ and an oracle $\mathcal{O}(q, v)$ outputting sets $(S_v)_{v \in c^{-1}(q)}$

```

1  $X(\mathbf{t}) := 0$ 
2 for  $q = 1$  to  $p$  do
3    $(S_v)_{v \in c^{-1}(q)} := \mathcal{O}(q, X)$  /* Use of the oracle */
4   foreach  $v \in c^{-1}(q)$  do  $X(v) := \max S_v$ 
5   repeat
6      $X_{pre} := X$ 
7     foreach  $v \in c^{-1}(q) \cap V_{\text{Max}}$  do
8        $X(v) := \max_{v' \in E(v)} (\omega(v, v') + X_{pre}(v'))$ 
9     foreach  $v \in c^{-1}(q) \cap V_{\text{Min}}$  do
10       $X(v) := \min_{v' \in E(v)} (\omega(v, v') + X_{pre}(v'))$ 
11     foreach  $v \in c^{-1}(q)$  do  $X(v) := \max (S_v \cap [-\infty, X(v)])$ 
12   until  $X = X_{pre}$ 
13 return  $X$ 

```

Algorithm 4: Accelerated value iteration algorithm for total-payoff games

Input: Total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP} \rangle$, SCC-decomposition $c: V \rightarrow \{0, 1, \dots, p\}$ and an oracle $\mathcal{O}(q, v)$ outputting sets $(S_v)_{v \in c^{-1}(q)}$

```

1 foreach  $v \in V$  do  $Y(v) := -\infty$ 
2 for  $q = 0$  to  $p$  do
3    $(S_v)_{v \in c^{-1}(q)} := \mathcal{O}(q, Y)$  /* Use of the oracle */
4   repeat
5     foreach  $v \in c^{-1}(q)$  do
6        $Y_{pre}(v) := Y(v); Y(v) := \max(0, Y(v)); X(v) := \max S_v$ 
7     repeat
8        $X_{pre} := X$ 
9       foreach  $v \in V_{\text{Max}} \cap c^{-1}(q)$  do
10         $X(v) := \max_{v' \in E(v)} [\omega(v, v') + \min(X_{pre}(v'), Y(v'))]$ 
11       foreach  $v \in V_{\text{Min}} \cap c^{-1}(q)$  do
12         $X(v) := \min_{v' \in E(v)} [\omega(v, v') + \min(X_{pre}(v'), Y(v'))]$ 
13       foreach  $v \in c^{-1}(q)$  do  $X(v) := \max (S_v \cap [-\infty, X(v)])$ 
14     until  $X = X_{pre}$ 
15      $Y := X$ 
16     foreach  $v \in V$  such that  $Y(v) > (|V| - 1)W$  do  $Y(v) := +\infty$ 
17   until  $Y = Y_{pre}$ 
18 return  $Y$ 

```

most L distinct weights (that can be arbitrarily large in absolute value) satisfies this criterion. Table 1 contains the results obtained with the heuristics on the parametric example presented before. Observe that the acceleration technique permits here to decrease drastically the execution time, the number of iterations in both loops depending not even anymore on W . Even though the number of iterations in the external loop increases with heuristics, due to the decomposition, less computation is required in each internal loop since we only apply the computation for the active component.

6 Conclusion

In this work, we have provided the first (to the best of our knowledge) pseudo-polynomial time algorithm to solve total-payoff games with arbitrary (positive and negative) weights. This algorithm is a variation on the classical value iteration technique. To obtain this algorithm, we have reduced the problem of solving total-payoff games to that of solving min-cost reachability (MCR) games, a variant of the former where the game stops as soon as a target vertex is reached (in which case the payoff of the plays is the total accumulated weight of the play up to the target). We believe that those MCR games are interesting by themselves, as they can be used to model problems where, for instance, a target configuration must be reached while ensuring a minimal energy spending. Notice also that they have been used as a building block for the resolution of priced timed games in Brihaye et al (2014, 2015a). We have characterised the optimal strategies that one can extract in those total-payoff and MCR games. Finally, we have implemented our algorithms and proposed some heuristics that take into account the structure of the games to speed up the computation. As future works, we would like to push further the MCR games in a context of non-zero sum games where each player wants to optimise its accumulated cost until reaching its own target. As a possible direction, the search for Nash equilibria in this context will most likely benefit from our better understanding of optimal strategies for both players in the underlying zero-sum games. This bridge from zero-sum to non-zero-sum games has already been investigated for concurrent priced games by Klimoš et al (2012), and by Brihaye et al (2013) to find simple Nash equilibria for large classes of multiplayer cost games.

Acknowledgements The authors are indebted to Jean-François Raskin for enlightening discussions, and, in particular, for suggesting the example in Fig. 2. They also want to thank reviewers of this version, as well as the short version Brihaye et al (2015b), that helped greatly improving its content.

References

- Björklund H, Vorobyov S (2007) A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics* 155:210–229
- Brihaye T, De Pril J, Schewe S (2013) Multiplayer cost games with simple nash equilibria. In: *Proceedings of the International Symposium on Logical Foundations of Computer Science (LFCS'13)*, Springer, Lecture Notes in Computer Science, vol 7734, pp 59–73
- Brihaye T, Geeraerts G, Krishna SN, Manasa L, Monmege B, Trivedi A (2014) Adding negative prices to priced timed games. In: *Proceedings of the 25th International Conference on Concurrency Theory (CONCUR'14)*, Springer, Lecture Notes in Computer Science, vol 8704, pp 560–575, DOI 10.1007/978-3-662-44584-6_38
- Brihaye T, Geeraerts G, Haddad A, Lefauchaux E, Monmege B (2015a) Simple priced timed games are not that simple. In: *Proceedings of the 35th Confer-*

- ence on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'15), Schloss Dagstuhl–Leibniz-Zentrum für Informatik, LIPIcs
- Brihaye T, Geeraerts G, Haddad A, Monmege B (2015b) To reach or not to reach? Efficient algorithms for total-payoff games. In: Proceedings of the 26th International Conference on Concurrency Theory (CONCUR'15), Schloss Dagstuhl–Leibniz-Zentrum für Informatik, LIPIcs, vol 42, pp 297–310
- Brim L, Chaloupka J, Doyen L, Gentilini R, Raskin JF (2011) Faster algorithms for mean-payoff games. *Formal Methods for System Design* 38(2):97–118
- Chen T, Forejt V, Kwiatkowska M, Parker D, Simaitis A (2013) Automatic verification of competitive stochastic systems. *Formal Methods in System Design* 43(1):61–92
- Comin C, Rizzi R (2015) An improved pseudo-polynomial upper bound for the value problem and optimal strategy synthesis in mean payoff games. Tech. Rep. 1503.04426, arXiv
- Ehrenfeucht A, Mycielski J (1979) Positional strategies for mean payoff games. *International Journal of Game Theory* 8(2):109–113
- Filiot E, Gentilini R, Raskin JF (2012) Quantitative languages defined by functional automata. In: Proceedings of the 23rd International Conference on Concurrency theory (CONCUR'12), Springer, Lecture Notes in Computer Science, vol 7454, pp 132–146
- Gawlitzka TM, Seidl H (2009) Games through nested fixpoints. In: Proceedings of the 21st International Conference on Computer Aided Verification (CAV'09), Springer, Lecture Notes in Computer Science, vol 5643, pp 291–305
- Gimbert H, Zielonka W (2004) When can you play positionally? In: Proceedings of the 29th International Conference on Mathematical Foundations of Computer Science (MFCS'04), Springer, Lecture Notes in Computer Science, vol 3153, pp 686–698
- Khachiyan L, Boros E, Borys K, Elbassioni K, Gurvich V, Rudolf G, Zhao J (2008) On short paths interdiction problems: Total and node-wise limited interdiction. *Theory of Computing Systems* 43:204–233
- Klimoš M, Larsen KG, Štefaňák F, Thaarup J (2012) Nash equilibria in concurrent priced games. In: Proceedings of the 6th international conference on Language and Automata Theory and Applications (LATA'12), Springer, Lecture Notes in Computer Science, vol 7183, pp 363–376
- Martin DA (1975) Borel determinacy. *Annals of Mathematics* 102(2):363–371
- Puterman ML (1994) *Markov Decision Processes*. John Wiley & Sons, Inc., New York, NY
- Strauch RE (1966) Negative dynamic programming. *The Annals of Mathematical Statistics* 37:871–890
- Tarjan RE (1972) Depth first search and linear graph algorithms. *SIAM Journal on Computing* 1(2):146–160
- Thomas W (1995) On the synthesis of strategies in infinite games. In: Symposium on Theoretical Aspects of Computer Science (STACS'95), Springer, Lecture Notes in Computer Science, vol 900, pp 1–13
- Thuijsman F, Vrieze OJ (1987) The bad match; a total reward stochastic game. *Operations-Research-Spektrum* 9(2):93–99, DOI 10.1007/BF01732644
- Zwick U, Paterson MS (1996) The complexity of mean payoff games. *Theoretical Computer Science* 158:343–359

A Comparison with the *longest shortest path problem* of Björklund and Vorobyov

Björklund and Vorobyov have studied games related to our min-cost reachability games, as the longest shortest path problem (LSP for short) in (Björklund and Vorobyov, 2007). To clarify the following discussion, let us recall the definition of LSP, adapted to our syntax.

The LSP problem considers a weighted graph $\mathcal{G}$, whose vertices are partitioned into two players, and equipped with a single target vertex $\mathbf{t}$. The problem asks to find a *memoryless* strategy σ_{Max} of Max such that in the graph $\mathcal{G}_{\sigma_{\text{Max}}}$, obtained from $\mathcal{G}$ by deleting all outgoing edges from vertices of V_{Max} except those selected in σ_{Max} , the length of the shortest path from every vertex to the target $\mathbf{t}$ is as large as possible (over all memoryless strategies). The definition of paths from a vertex v to the target is, however, very different from ours: such a path may indeed be a finite play, without cycle, from v to the target $\mathbf{t}$ (in which case its length is the sum of the weights of edges). However, every cycle containing v (even if it can not reach the target) is also considered as a path to the target. The *length* of such cycle is defined as follows:

1. if the cycle has a negative weight, its length is $-\infty$;
2. else, if it has a positive weight, its length is $+\infty$;
3. else, if it has zero weight, its length is 0.

Consider, as an example, the weighted graph of Fig. 9, where, once again, Max vertices are depicted with circles. Clearly, vertices $\{v_1, v_2, v_3\}$ and $\{v_4\}$ form independant subgames. We first consider the status of v_1 , v_2 and v_3 . In this subgame, Max has two possible strategies: σ_1 that selects the (v_1, v_2) edge, or σ_2 that selects the (v_1, v_3) edge instead.

1. If we select σ_1 , we obtain the following values for v_1 , v_2 and v_3 respectively: 0, 0 and 1. Indeed, from vertex v_2 , we can loop in the zero cycle in-between v_1 and v_2 , which is better than the simple path to the target (having weight 3), so that the distance from v_1 and v_2 is 0. From vertex v_3 , only a single path leads to the target with weight 1.
2. If we select σ_2 , we obtain the following values for v_1 , v_2 and v_3 respectively: 1, 3 and 1. Indeed, from vertex v_3 , the distance is the length of the simple path to the target, which is better than looping in-between v_1 and v_3 with a positive weight cycle. Hence, from vertex v_1 , the distance is 1 too, whereas from v_2 , the distance is 3.

Now, let us turn our attention to v_4 . Again, we must consider two strategies: strategy σ_3 selects the (v_4, v_4) edge (the self-loop on v_4) and strategy σ_4 selects the $(v_4, \mathbf{t})$ edge.

1. If we select σ_3 , we obtain value $-\infty$ for v_4 , because all plays starting in v_4 loop in this negatively priced cycle. Recall that with our definition of min-cost reachability, such a cycle that never reaches the target would yield a total-payoff of $+\infty$, regardless of the weights. This is coherent with our intuition that the first goal of Min is to reach the target.
2. If we select σ_4 , the value of v_4 is now 0, which is thus better for Max.

Finally, the optimal strategy for Max is to choose edges (v_1, v_3) and $(v_4, \mathbf{t})$, leading to the LSP distances $(1, 3, 1, 0)$ for vertices (v_1, v_2, v_3, v_4) . With our definition of min-cost reachability game, the situation is *completely different*: the value vector is then $(2, 3, 1, +\infty)$, associated with the optimal strategy σ_{Max} selecting edges (v_1, v_2) and (v_4, v_4) .

Indeed, two main differences separate our two definitions. The first one is the treatment of negative weight cycles. Actually, in LSP, the fact that after selecting strategy σ_{Max} a vertex can not reach the target anymore in $\mathcal{G}_{\sigma_{\text{Max}}}$ does not prevent from mapping the distance $-\infty$ to a vertex contained in a negative cycle. This is in contrast with our definition, that would benefit to Max since in such a situation, the target is not reachable leading to the value $+\infty$.⁹ The second difference consists in the treatment of zero weight cycle. In LSP, a distance zero is then computed, which is highly related with the fact that the authors want to apply the resolution of LSP to mean-payoff games.

Nevertheless, in Section 9 of (Björklund and Vorobyov, 2007), the authors *briefly* study another more natural definition, mapping zero weight cycle to the distance $+\infty$ (closer to our definition). In that case, the LSP distances of vertices v_1 , v_2 and v_3 then match our value vector in min-cost reachability games. Unfortunately, it is not clear how the algorithm that is presented by Björklund and Vorobyov can be adapted to accomodate this new definition

⁹ We believe that this difference would certainly be eliminated by our preprocessing of vertices of value $+\infty$ presented in the first item of Theorem 3.

Fig. 9 An instance of the longest shortest path problem

of the cost and compute the values of the nodes. Indeed, Proposition 9.1 of (Björklund and Vorobyov, 2007) proves that the decision version of the LSP (i.e. deciding if the LSP of a vertex is positive or not in a given graph) is in $\text{NP} \cap \text{co-NP}$, and they claim (without formal proof) that their pseudo-polynomial time algorithm permits to decide this problem. This requires a first transformation of the problem, so that no zero weight cycles remain in the weighted graph. This transformation, explained above Proposition 9.1, is correct but *does not preserve the LSP of the vertices*, so that their algorithm do not compute the LSP distance of vertices, but only study their positivity.

Hence, the comparison between our definition of min-cost reachability and the definition of LSP can be summarised as follows:

1. Either one relies on the first definition of LSP given in the article. In this case, Björklund and Vorobyov propose a pseudo-polynomial time algorithm to compute the LSP, but these values *do not match our definition of the min-cost reachability* payoff, as demonstrated by the above example.
2. Or one relies on the second definition. In this case, we believe that the LSP values may match our definition of min-cost reachability but the article *does not explain formally how to compute those new values*.

It is plausible that the pseudo-polynomial time algorithm of Björklund and Vorobyov could be adapted to compute the LSP value according to the second definition. Yet, even in this case, there are several points of comparison worth mentioning:

1. The worst-case complexity of the algorithm proposed by Björklund and Vorobyov is $O(|V|^2|E|W)$, which matches ours (see Proposition 5). Nevertheless, our solution is much easier to describe and to implement: while they must make repeated calls to a modified version of the Bellman-Ford algorithm (the modification is crucial to obtain their complexity), we have a simple fixed point algorithm.
2. Our algorithm exploits the value iteration paradigm (that can be seen as a *backward induction*), while theirs is a *strategy iteration algorithm*. Because of that, there are examples on which our algorithm is more efficient than theirs (although the worst-case complexity is the same). As an example, the LSP instance presented in Fig. 2 of (Björklund and Vorobyov, 2007), with $2n + 1$ vertices but a biggest weight W exponential in n , requires an exponential number 2^n of iterations for the strategy iteration algorithm, but our value iteration algorithm (based on backward induction) would solve it in linear time with respect to n . More precisely, with our tool, we are able to compute the values of this game in less than a millisecond for constant $n = 15$ (i.e. a game with 32 vertices and largest weight $W \approx 2.68 \cdot 10^8$): the value of the initial state obtained is $4.74 \cdot 10^9$, and as aforementioned, the number of iterations that our value iteration requires is 16, linear in n . Our tool being an add-on of the PRISM model-checker, relying on the use of integer rewards, we have not been able to build the game for a value n greater than 15.
3. We propose several acceleration heuristics that perform well on several examples we have tried. These accelerations cannot easily be incorporated in their algorithm (because it is a strategy iteration algorithm).
4. Finally, from the theoretical point of view, Björklund and Vorobyov do not study the strategies of the opponent player Min. In contrast, our study allows us to produce optimal strategies for both players, and in particular, show that memoryless strategies may not be sufficient for Min, whereas they are enough for Max.