The Influence of Neighbourhood and Choice on the Complexity of Finding Pure Nash Equilibria

Felix Fischer [a,b] Markus Holzer [b] Stefan Katzenbeisser [b]

[a] *Institut für Informatik, Universität München, Oettingenstraße 67, 80538 München, Germany*

[b] *Institut für Informatik, Technische Universität München, Boltzmannstraße 3, 85748 Garching, Germany*

Key words: Computational complexity, Game theory, Pure strategy Nash equilibrium

1 Introduction

Game theory is a mathematical framework for representing interactions of rational players trying to achieve possibly contradictory goals. Players' strategies are said to be in *Nash equilibrium* if no player can do better by unilaterally changing her strategy. Nash's famous theorem [1] guarantees the existence of at least one such equilibrium when players are allowed to play *mixed strategies*, i.e., probabilistic combinations of actions. If strategies are restricted to deterministic choice of actions, called *pure strategies*, the existence of an equilibrium is no longer guaranteed (see, e.g., [2]). Equilibria in the latter case are referred to as pure strategy Nash equilibria, or pure Nash equilibria in short. Complexity issues related to pure Nash equilibria have recently been investigated in [3]. It was shown that even for a very restricted class of games in graphical normal form, where each player is allowed to play at most 3 different actions and her payoff depends on at most 3 other players, determining whether a given game has a pure Nash equilibrium is NP-complete. Moreover, some tractable classes of strategic games have been identified.

In this paper, we strengthen the above mentioned NP-completeness result to apply to an even more restrictive setting. To be precise, we show that 2 actions per player, 2-bounded neighbourhood, and 2-valued payoff matrices suffice to prove NP-completeness. It should be noted that this is the best possible result, because

Email addresses: `fischerf@tcs.ifi.lmu.de` (Felix Fischer), `holzer@in.tum.de` (Markus Holzer), `katzenbe@in.tum.de` (Stefan Katzenbeisser).

deciding the existence of a pure Nash equilibrium becomes trivial in the case of a single action for each player and tractable for 1-bounded neighbourhood. In fact, we show the latter problem to be NL-complete in general, and thus solvable in deterministic polynomial time. Interestingly, it turns out that the number of actions in a game with 1-bounded neighbourhood is a sensitive parameter: Restricting the number of actions for each player to a constant makes the problem even easier than NL unless L = NL. In this way, we obtain a nice alternative characterisation of the deterministic-nondeterministic problem for logspace bounded Turing machines in terms of the number of actions for strategic games with 1-bounded neighbourhood.

2 Preliminaries

We recall some preliminaries on strategic games and refer to [2] for a comprehensive introduction.

Definition 1 (Strategic game) *A* game in strategic form *is a* 4*-tuple* $\mathscr{G} = (P, (A_p)_{p\in P}, (N_p)_{p\in P}, (u_p)_{p\in P})$*, where* P *is a set of* players *and for each player* $p \in P$*,* A_p *is a nonempty set of* actions *available to* p*,* $N_p \subseteq P \setminus \{p\}$ *is a set of* neighbours *of* p*, and* $u_p : (\times_{q\in N_p\cup\{p\}} A_q) \to \mathbb{Q}$ *is a function mapping each combination of actions for* p *and her neighbours to a rational* payoff *for* p*.*

This definition of a strategic game differs from that commonly used in game theory in that it explicitly represents the neighbourhood of a player as the set of other players that potentially have an influence on her payoff. A game is said to have *k-bounded neighbourhood* if there is a constant k such that $|N_p| \leq k$ for each $p \in P$. For any k, games with k-bounded neighbourhood form a subclass of strategic games.

For a particular player p, her choice to play action $a \in A_p$ deterministically is also referred to as a *pure strategy*. For a non-empty set $P' \subseteq P$ of players, an *action profile* for P' is a vector containing exactly one action for each player in P'. The different action profiles for P correspond to the possible outcomes of game $\mathscr{G}$. Players' payoff functions u_p naturally extend to $\times_{p\in P} A_p$, and we will denote these extensions by u_p as well. A *pure (strategy) Nash equilibrium* for $\mathscr{G}$ is an action profile where no player can increase her own payoff by unilaterally changing her action.

Definition 2 (Pure strategy Nash equilibrium) *An action profile* $a \in \times_{p\in P} A_p$ *is called a* pure strategy Nash equilibrium *if and only if for each player* $p \in P$ *and each action* $a'_p \in A_p$*,*

$$u_p(a) \geq u_p((a_{-p}, a'_p)),$$

where (a_{-p}, a'_p) *denotes the action profile where player* p *plays* a'_p *and all other players play the same action as in* a*.*

If the number of players is small, deciding whether a game has a pure Nash equilibrium can simply be done by enumerating all possible action profiles and checking if one of them satisfies the equilibrium condition. The problem becomes interesting, though, for games that admit a compact representation when the number of players is growing. Clearly, games with k-bounded neighbourhood belong to this class.

We further assume the reader to be familiar with the basics of complexity theory (see, e.g., [4]). In particular we consider the well known chain of complexity classes $\mathrm{L} \subseteq \mathrm{NL} \subseteq \mathrm{P} \subseteq \mathrm{NP}$. Here L (NL, respectively) denotes the class of all languages accepted by deterministic (nondeterministic, respectively) logarithmic space bounded Turing machines, and P (NP, respectively) refers to the class of all languages accepted by deterministic (nondeterministic, respectively) polytime bounded Turing machines. Moreover, #P is the class of functions $f(x)$ for which there is a polytime nondeterministic Turing machine M such that $f(x)$ equals the number of accepting computations of M on x. The class #L is defined analogously for nondeterministic logspace Turing machines.

All problems investigated in this paper take strategic games as inputs. We therefore need an appropriate coding function $\langle\cdot\rangle$ mapping a strategic game $\mathscr{G}$ to a word $\langle\mathscr{G}\rangle$ over a fixed alphabet. We will not go into the details of the definition of $\langle\cdot\rangle$, but assume that it satisfies standard properties. In particular, the payoff functions are explicitly represented by a single table or matrix having an entry for every action profile a, displaying a list of payoffs $u_p(a)$ for each player $p \in P$. If the payoff function for each player p is represented by a separate table or matrix containing an entry $u_p(a)$ for every action profile $a \in \times_{q \in N_p \cup \{p\}} A_p$, then the game is said to be in *graphical normal form* [5].

3 Hard and Easy Games

The following theorem states that deciding the existence of a pure Nash equilibrium is hard even if each player has a 2-bounded neighbourhood, 2-valued payoff matrices and only 2 different actions. We improve upon a recent result by [3].

Theorem 3 *Deciding whether a strategic game $\mathscr{G}$ has a pure strategy Nash equilibrium is NP-complete. Hardness holds even if $\mathscr{G}$ is in graphical normal form, has* 2*-bounded neighbourhood and* 2*-valued payoff matrices, and where each player is allowed to play at most* 2 *different actions.*

PROOF. *Membership* in NP is obvious. We can guess an action profile a, and verify in polynomial time whether a is a pure Nash equilibrium — for the latter task we have to check whether $u_p((a_{-p}, a'_p)) \leq u_p(a)$ for each player $p \in P$ and for each action $a'_p \in A_p$.

For *hardness* we argue as follows: Recall that circuit satisfiability, i.e., deciding whether for a given Boolean circuit $\mathscr{C}$ with n inputs and one output there exists an assignment such that $\mathscr{C}$ evaluates to *true*, is NP-complete (see, e.g., [4]). W.l.o.g, assume that $\mathscr{C}$ contains at least one input and one (internal) gate, and that *NOT* gates only occur at the input layer; for an arbitrary circuit, all *NOT* gates can be moved to the input layer in polynomial time by successive application of de Morgan's law.

We define a game $\mathscr{G}$ in graphical normal form for a Boolean circuit $\mathscr{C}$, providing a polynomial-time reduction of satisfiability of $\mathscr{C}$ to the problem of finding a pure Nash equilibrium in $\mathscr{G}$. As for the players of G, there is one for each circuit input, one for each (positive or negative) literal, and one for each gate of the circuit (of types *AND* and *OR*). We write P_i, P_x, $P_{\overline{x}}$, $P_\wedge$, and $P_\vee$ to denote the respective sets of players. The output of the circuit corresponds to a particular player $o \in P_\wedge \cup P_\vee$. For each input player $p \in P_i$, let N_p be the empty set. For each gate player $g \in P_\wedge \cup P_\vee$, let N_g be the set of players corresponding to the inputs of the gate — whenever a gate is connected to one of the n inputs or to a *NOT*-gate, N_g contains the player corresponding to the appropriate positive or negative literal. Finally, for each literal player $\ell \in P_x \cup P_{\overline{x}}$, let N_ℓ contain the appropriate input player $i \in P_i$ and the output player o. Let the set of possible actions equal $\{t, f\}$ for every player $p \in P$, in which t and f can be interpreted as truth values. For an action profile a, let the payoff functions be defined as follows:

- For each input player $i \in P_i$, her payoff function u_i is such that $u_i(a) = 1$ regardless of the selected action.
- For each positive literal player $\ell \in P_x$, her payoff function u_ℓ is such that
  - $u_\ell(a) = 1$, if (i) ℓ plays the same action as the input player $i \in N_\ell$ and the output player o plays t, or (ii) ℓ plays t and o plays f;
  - $u_\ell(a) = 0$ in all other cases.
- For each negative literal player $\ell \in P_{\overline{x}}$, her payoff function u_ℓ is such that
  - $u_\ell(a) = 1$, if (i) ℓ plays the opposite of the action of the input player $i \in N_\ell$ and the output player o plays t, or (ii) ℓ plays t and o plays f;
  - $u_\ell(a) = 0$ in all other cases.
- For each gate player $g \in P_\wedge$, her payoff function u_g is such that
  - $u_g(a) = 1$, if (i) g and both players $a, b \in N_g$ play t, or (ii) g plays f and at least one of a and b plays f;
  - $u_g(a) = 0$ in all other cases.
- For each gate player $g \in P_\vee$, her payoff function u_g is such that
  - $u_g(a) = 1$, if (i) g plays t and at least one player $a, b \in N_g$ plays t, or (ii) g and both a and b play f;
  - $u_g(a) = 0$ in all other cases.

Payoffs for all types of players are summarised in Tables 1 and 2. We claim that $\mathscr{C}$ is satisfiable if and only if $\mathscr{G}$ has a pure Nash equilibrium.

		u_i	
$i \in P_i$	t		1
	f		1

			$i,o \in N_x$			
		u_x	tt	tf	ft	ff
$x \in P_x$	t		1	1	0	1
	f		0	0	1	0

			$i,o \in N_x$			
		$u_{\overline{x}}$	tt	tf	ft	ff
$x \in P_{\overline{x}}$	t		0	1	1	1
	f		1	0	0	0

Table 1
Payoffs for "input", "positive literal", and "negative literal" players, the latter two depending on the actions of the corresponding input player i and the output player o.

			$a,b \in N_g$			
		$u_\wedge$	tt	tf	ft	ff
$g \in P_\wedge$	t		1	0	0	0
	f		0	1	1	1

			$a,b \in N_g$			
		$u_\vee$	tt	tf	ft	ff
$g \in P_\vee$	t		1	1	1	0
	f		0	0	0	1

Table 2
Payoffs for "gate" players, depending on the actions of players a and b corresponding to the inputs of the gate.

The implication from left to right is seen as follows: Assume that $\mathscr{C}$ is satisfiable and consider a particular satisfying assignment ϕ. Consider further an action profile a for $\mathscr{G}$ where

- each input player plays according to ϕ,
- each literal player correctly reproduces the action of the corresponding input player (positive literals playing the same action as their input, negative ones playing the opposite action), and
- each gate player correctly implements the truth function of the respective gate depending on the inputs, i.e., actions of her neighbours.

Observe that following the definition of $\mathscr{G}$, and since ϕ is a satisfying assignment, player o will play t. In this case each player receives a payoff of 1, and since 1 is the maximum payoff, a is a Nash equilibrium for $\mathscr{G}$.

Conversely, it remains to be shown that every pure Nash equilibrium of $\mathscr{G}$ corresponds to a satisfying assignment of $\mathscr{C}$. We exploit the following properties of action profiles for $\mathscr{G}$:

P_1: A profile where all "literal" players $\ell \in P_x \cup P_{\overline{x}}$ play t and o plays f cannot be a pure Nash equilibrium. Since no negations occur above the literal players, any gate player $g \in P_\wedge \cup P_\vee$ with $N_g \subseteq (P_x \cup P_{\overline{x}})$ who is not playing t can improve her payoff by switching to t. By induction over the structure of the gate this holds for all gate players $g \in P_\wedge \cup P_\vee$, and for o in particular. This is a contradiction.

P_2: A profile where o plays f cannot be a pure Nash equilibrium. If o plays f, any literal player not playing t can improve her payoff by playing t. This contradicts P_1.

P_3: A profile where o plays t and any literal player $\ell \in P_x \cup P_{\overline{x}}$ plays an action that does not correctly implement the value of the corresponding input cannot be a Nash equilibrium. In this case, ℓ could change her action to increase her payoff, which is a contradiction.

P_4: A profile in which some gate player $g \in P_\wedge \cup P_\vee$ plays an action that does not correctly implement the Boolean function of the corresponding gate cannot be a Nash equilibrium. In this case, g could change her action to increase her payoff, which is a contradiction.

P_5: A profile where the input players do not play a satisfying assignment but o plays t cannot be a Nash equilibrium. Assuming such an equilibrium, and due to P_4, all gate players would have to play according to the Boolean function they implement. Thus, because the assignment is not satisfying, player o could change her action to f to increase her payoff. This is a contradiction.

By combining properties P_1 to P_5, we conclude that the only action profiles for $\mathscr{G}$ that are pure Nash equilibria are those corresponding to satisfying assignments of the Boolean circuit $\mathscr{C}$. Thus, there is a one-to-one correspondence between satisfying assignments of $\mathscr{C}$ and pure Nash equilibria of $\mathscr{G}$.

Further observe that the payoff matrices in Tables 1 and 2 can be constructed from a particular circuit in polynomial time, that each player in $\mathscr{G}$ has at most two neighbours and two different actions, and that there are only two distinct payoff values. □

This proof also shows the following corollary.

Corollary 4 *Computing the number of pure Nash equilibria of a strategic game $\mathscr{G}$ is #P-complete, even if $\mathscr{G}$ is in graphical normal form, has 2-bounded neighbourhood and 2-valued payoff matrices, and each player is allowed to play at most 2 different actions.* □

Obviously, when reducing the number of different actions per player to a single one, or considering single-valued payoff matrices, deciding whether a game has a pure Nash equilibrium becomes trivial. Hence, the only interesting case that remains are games with 1-bounded neighbourhood. The interaction among players of a strategic game $\mathscr{G}$ can be represented as a directed *interaction graph* $G(\mathscr{G}) = (P,E)$, where the vertices correspond to the players of $\mathscr{G}$, and $(p,q) \in E$ if and only if q is a neighbour of p, i.e., $q \in N_p$. Observe that for a game $\mathscr{G}$ with k-bounded neighbourhood, each node of $G(\mathscr{G})$ has an outdegree of at most k. It is thus easy to see that for a strategic game $\mathscr{G}$ with 1-bounded neighbourhood, $G(\mathscr{G})$ is a forest of "trycles" (we use this neologism to denote trees to which at most *one* cycle has been attached). Formally, such a *trycle* is either (1) a rooted tree with edge orientation toward the root or (2) a rooted tree with edge orientation toward the root and an additional edge from the designated root to some *other* node in the tree. The following lemma states that the existence of a pure Nash equilibrium in a strategic game whose interaction graph is a trycle only depends on the interaction of the players in the cycle.

We omit the straightforward proof.

Lemma 5 *Let $\mathscr{G} = (P,(N_p)_{p\in P},(A_p)_{p\in P},(u_p)_{p\in P})$ be a strategic game such that $G(\mathscr{G})$ is a* trycle. *Then, $\mathscr{G}$ has a pure strategy Nash equilibrium if and only if the game $\mathscr{G}'$ obtained from $\mathscr{G}$ by restricting the set of players to those in the unique cycle of $G(\mathscr{G})$, i.e., $\mathscr{G}' = (P',(N_p)_{p\in P'},(A_p)_{p\in P'},(u_p)_{p\in P'})$ with $P' = \{\, p \in P \mid \text{there is a path from } p \text{ to } p \text{ in } G(\mathscr{G}) \,\}$, has a pure strategy Nash equilibrium.* □

We are now ready to classify the complexity of the pure Nash equilibrium problem for strategic games with 1-bounded neighbourhood.

Theorem 6 *Deciding whether a strategic game $\mathscr{G}$ with* 1*-bounded neighbourhood has a pure strategy Nash equilibrium is NL-complete. Hardness holds even if $\mathscr{G}$ is in graphical normal form and has* 2*-valued payoff matrices.*

PROOF. *Membership* in NL is seen as follows: Let $\mathscr{G}$ be a strategic game with 1-bounded neighbourhood. By combining the above observation on $G(\mathscr{G})$ and Lemma 5, it is sufficient to reduce each trycle in the forest to a (possibly empty) cycle. This can be done in deterministic logspace, since each vertex of the interaction graph $G(\mathscr{G})$ has outdegree at most 1. We thus end up with a forest of cycles. We can now guess an action profile for each cycle independently in the following way. We start at an arbitrary (deterministically computed) node in the cycle and guess an action for the player corresponding to this node. We then traverse the cycle backwards, guessing an action for each player and checking whether this action is the best response w.r.t. the action of the next player. The computation ends when the initially chosen node has been reached, accepting if and only if the initially chosen action equals the action guessed in the last step of the traversal. Note that the backwards traversal can be done in deterministic logarithmic space; in addition, we only need to maintain a constant number of pointers. Thus, the whole process can be performed on a logspace bounded Turing machine. If for each cycle an individual pure Nash equilibrium exists, then the original strategic game $\mathscr{G}$ also has a pure Nash equilibrium.

It remains to show NL-*hardness*. For this, we reduce the s-t-reachability problem in directed graphs, i.e., deciding whether for a given directed graph $G = (V,E)$, with $E \subseteq V \times V$, and two designated vertices $s,t \in V$ there exists a path from s to t in G (see, e.g., [4]), to finding a pure Nash equilibrium in a game with 1-bounded neighbourhood. W.l.o.g., assume that (1) $V = \{1,2,\ldots,n\}$ with $n \geq 4$, (2) $s = 1$ and $t = n$, (3) for each node i with $1 \leq i \leq n-1$ there exists at least one edge leaving i, and (4) the *only* edge leaving n is a self-loop.

We then define a game $\mathscr{G}$ in graphical normal form for a directed graph G. The set of players of $\mathscr{G}$ is given by $P = \{p_1, p_2, \ldots, p_{n-1}\}$, neighbourhood is given by $N_{p_1} = \{p_{n-1}\}$ and $N_{p_t} = \{p_{t-1}\}$ for $2 \leq t \leq n-1$. The set of possible actions, which is the same for all players, is given by $\{\, a_{ij} \mid 1 \leq i, j \leq n \,\}$, in which a_{ij} can be interpreted as selection of (the possibly non-existing) edge (i, j) in graph G. Let a denote an action profile, then the payoff functions are defined as follows:

- Player p_1's payoff function u_1 is such that
  - $u_1(a) = 1$ if (i) p_1 plays a_{1i} for some i with $(1,i) \in E$ and p_{n-1} plays a_{jn} for some j, or (ii) p_1 plays a_{nn} and p_{n-1} plays an action a_{jk} such that $k \neq n$;
  - $u_1(a) = 0$ in all other cases.
- For each player $p_t \in P$ with $2 \leq t \leq n-1$, her payoff function u_t closely resembles the transition matrix of G; more precisely,
  - $u_t(a) = 1$ if p_t plays a_{jk} for some j with $(j,k) \in E$ and p_{t-1} plays a_{ij} for some i;
  - $u_t(a) = 0$ in all other cases.

We claim that there is a path from node 1 to node n in graph G if and only if game $\mathscr{G}$ has a pure Nash equilibrium.

The implication from left to right is seen as follows: Assume that in G there exists a path $1 = i_1, i_2, \ldots, i_n = n$ of length n (a shorter path can be extended by virtue of the self-loop at node n). Consider further an action profile a for $\mathscr{G}$ where each player p_t, with $1 \leq t \leq n-1$, plays action $a_{i_t i_{t+1}}$. In this case, each player receives a payoff of 1, and since 1 is the maximum payoff, a is a Nash equilibrium of $\mathscr{G}$.

Conversely, it remains to be shown that every Nash equilibrium for $\mathscr{G}$ corresponds to a path linking nodes 1 and n in G. We exploit the following properties of action profiles for $\mathscr{G}$:

P_1: A profile where player p_t with $2 \leq t \leq n-1$ plays $a_{k\ell}$ and p_{t-1} plays a_{ij} for some $j \neq k$ cannot be a Nash equilibrium. In this case, p_t would get a payoff of 0, and by construction there exists an alternative action $a_{k\ell}$ with $(k,\ell) \in E$ with payoff 1 (because every node has outdegree at least 1). Hence, p_t could change her action to increase her payoff, which is a contradiction.

P_2: A profile where p_t with $1 \leq t \leq n-2$ plays a_{in} for some $1 \leq i \leq n$ and some p_j with $t+1 \leq j \leq n-1$ does *not* play a_{nn} cannot be a Nash equilibrium. W.l.o.g., assume that j is the smallest such number. In this case, p_j could improve her payoff by playing a_{nn}, which is a contradiction.

P_3: A profile where p_{n-1} plays $a_{k\ell}$ with $\ell \neq n$ cannot be a pure Nash equilibrium. In this case, p_1 could improve her payoff by playing a_{nn}. This contradicts P_2.

P_4: A profile where p_1 plays a_{ij} with $i \neq 1$ cannot be a pure Nash equilibrium. We distinguish two different cases. In case p_{n-1} plays a_{kn} for some k, p_1 could improve her payoff by playing a_{1j}, which is a contradiction. If instead p_{n-1} plays $a_{k\ell}$ for some $\ell \neq n$, p_1 could improve her payoff by playing a_{nn}. This contradicts P_2.

By combining properties P_1 to P_4, we conclude that the only action profiles that are pure Nash equilibria for $\mathscr{G}$ are those where

- p_1 plays an action a_{1i} with $(1,i) \in E$,

- each pair p_{t-1} and p_t of players with $2 \leq t \leq n-2$ play actions a_{ij} and a_{jk}, respectively, for some $1 \leq i, j, k \leq n$, with $(j,k) \in E$, and
- player p_{n-1} plays an action a_{kn} with $(k,n) \in E$.

Thus, there is a one-to-one correspondence of paths from node 1 to n in G and pure Nash equilibria of $\mathscr{G}$.

Further observe that the payoff matrices can be constructed from a particular graph G in logarithmic space, that each player has at most one neighbour, and that there are only two distinct payoff values. □

An immediate consequence of the above proof is given next.

Corollary 7 *Computing the number of pure Nash equilibria of a strategic game with* 1*-bounded neighbourhood is #L-complete, even if the game is in graphical normal form.* □

Observe that in the hardness part of the proof of Theorem 6, the number of actions of game $\mathscr{G}$ is polynomially bounded in the number of nodes of graph G. If the number of actions is bounded by a constant, the complexity of deciding whether $\mathscr{G}$ has a pure Nash equilibrium can be improved to L-completeness.

Theorem 8 *Deciding whether a strategic game $\mathscr{G}$ with* 1*-bounded neighbourhood and a* constant number of actions *has a pure strategy Nash equilibrium is L-complete under NC^1 reductions. Hardness holds even if $\mathscr{G}$ is in graphical normal form and has* 2*-valued payoff matrices.*

PROOF. The proof for *membership* in L follows similar lines as the membership proof for Theorem 6, with the following difference: Because the number of actions is a constant, deterministic logspace is enough to determine whether a strategic game whose interaction graph is a (forest of) cycle(s) has a pure Nash equilibrium; a Turing machine can write down all best response actions for a particular player by running backwards through the cycle and updating this information step by step.

Hardness can be shown by means of a straightforward reduction from the *s*-*t*-reachability problem for directed graphs of outdegree one to the problem under consideration. We define a game $\mathscr{G}$ for a graph G of this kind, where now each player in $\mathscr{G}$ corresponds to a node in G (for simplicity, we use names for players and nodes interchangeably), neighbourhood in $\mathscr{G}$ is given by neighbourhood in G, and the actions are from $\{t, f\}$. Additionally, player t has player s as a neighbour. Let a be an action profile for $\mathscr{G}$, then the payoff functions are defined as follows:

- Player t's payoff function is such that
  - $u_t(a) = 1$ if t plays the *opposite* action of s;
  - $u_t(a) = 0$ in all other cases.

- The payoff function of any player $p \neq t$ who has a neighbour $q \neq t$ is such that
  - · $u_p(a) = 1$ if p plays the same as q;
  - · $u_p(a) = 0$ in all other cases.
- For any other player, the payoff function is such that $u_p(a) = 1$ regardless of the action.

It is easily verified that this reduction can be computed in NC^1. Further, $\mathscr{G}$ does *not* have a pure Nash equilibrium if and only if t is reachable from s in G (since both nodes are in the same tree and, by construction, on a cycle). Since deterministic logspace is closed under complement, this completes the proof. □

The proof even shows that the problem remains L-complete if the number of actions per player is a slowly growing function.

Corollary 9 *Deciding whether a strategic game $\mathscr{G}$ with 1-bounded neighbourhood and $\log\log n + O(1)$ actions has a pure strategy Nash equilibrium is L-complete under NC^1 reductions. Hardness holds even if $\mathscr{G}$ is in graphical normal form and has 2-valued payoff matrices.*

PROOF. To see this, observe that in the logspace algorithm described in the proof of Theorem 8, we can still write down all possible best responses within the space bound. □

By combining Theorems 6 and 8, we obtain an alternative characterisation of the deterministic-nondeterministic problem for logspace bounded Turing machines.

Corollary 10 *The following statements are equivalent:*

(1) $L = NL$.
(2) Deciding whether a strategic game $\mathscr{G}$ with 1-bounded neighbourhood has a pure strategy Nash equilibrium can be done in deterministic logarithmic space.
(3) For each strategic game $\mathscr{G}$ with 1-bounded neighbourhood, we can construct in deterministic logarithmic space a strategic game $\mathscr{G}'$ with 1-bounded neighbourhood and a constant *number of actions such that $\mathscr{G}$ has a pure strategy Nash equilibrium if and only if $\mathscr{G}'$ has a pure strategy Nash equilibrium.* □

Acknowledgements

We would like to thank Hermann Gruber for taking part in some fruitful discussions on the subject and the anonymous reviewers for their helpful comments.

References

[1] J. F. Nash, Non-cooperative games, Annals of Mathematics 54 (2) (1951) 286–295.

[2] R. B. Myerson, Game Theory: Analysis of Conflict, Harvard University Press, 1997.

[3] G. Gottlob, G. Greco, F. Scarcello, Pure Nash equilibria: Hard and easy games, Journal of Artificial Intelligence Research 24 (2005) 357–406.

[4] C. H. Papadimitriou, Computational Complexity, Addison-Wesley, 1994.

[5] M. Kearns, M. L. Littman, S. Singh, Graphical models for game theory, in: Proceedings of the 17th International Conference on Uncertainty in Artificial Intelligence (UAI), 2001, pp. 253–260.