Games with Secure Equilibria*,**

Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdziński

Department of Electrical Engineering and Computer Sciences,
University of California, Berkeley, USA
{c_krish, tah, mju}@eecs.berkeley.edu

Abstract. In 2-player non-zero-sum games, Nash equilibria capture the options for rational behavior if each player attempts to maximize her payoff. In contrast to classical game theory, we consider lexicographic objectives: first, each player tries to maximize her own payoff, and then, the player tries to minimize the opponent's payoff. Such objectives arise naturally in the verification of systems with multiple components. There, instead of proving that each component satisfies its specification no matter how the other components behave, it often suffices to prove that each component satisfies its specification provided that the other components satisfy their specifications. We say that a Nash equilibrium is *secure* if it is an equilibrium with respect to the lexicographic objectives of both players. We prove that in graph games with Borel winning conditions, which include the games that arise in verification, there may be several Nash equilibria, but there is always a unique maximal payoff profile of a secure equilibrium. We show how this equilibrium can be computed in the case of ω-regular winning conditions, and we characterize the memory requirements of strategies that achieve the equilibrium.

1 Introduction

We consider 2-player non-zero-sum games, i.e., non-strictly competitive games. A possible behavior of the two players is captured by a strategy profile (σ, π), where σ is a strategy of player 1, and π is a strategy of player 2. Classically, the behavior (σ, π) is considered *rational* if the strategy profile is a Nash equilibrium [7] — that is, if neither player can increase her payoff by unilaterally changing her strategy. Formally, let $v_1^{\sigma,\pi}$ be the real-valued payoff of player 1 if the strategies (σ, π) are played, and let $v_2^{\sigma,\pi}$ be the corresponding payoff of player 2. Then (σ, π) is a Nash equilibrium if (1) $v_1^{\sigma,\pi} \geq v_1^{\sigma',\pi}$ for all player 1 strategies σ', and (2) $v_2^{\sigma,\pi} \geq v_2^{\sigma,\pi'}$ for all player 2 strategies π'. Nash equilibria formalize a notion of rationality which is strictly *internal*: each player cares about her own payoff but does not in the least care (cooperatively or adversarially) about the other player's payoff.

* This research was supported in part by the ONR grant N00014-02-1-0671, the AFOSR MURI grant F49620-00-1-0327, and the NSF grant CCR-0225610.

** This is an extended version of the paper "Games with Secure Equilibria" that appeared in the proceedings of *Logic in Computer Science* (LICS), 2004.

F.S. de Boer et al. (Eds.): FMCO 2004, LNCS 3657, pp. 141–161, 2005.

Choosing Among Nash Equilibria. A classical problem is that many games have *multiple* Nash equilibria, and some of them may be preferable to others. For example, one might partially order the equilibria by $(\sigma, \pi) \succeq (\sigma', \pi')$ if both $v_1^{\sigma,\pi} \geq v_1^{\sigma',\pi'}$ and $v_2^{\sigma,\pi} \geq v_2^{\sigma',\pi'}$. If a *unique maximal* Nash equilibrium exists in this order, then it is preferable for both players. However, maximal Nash equilibria may not be unique. In these cases *external* criteria, such as the sum of the payoffs for both players, have been used to evaluate different rational behaviors [9,14]. These external criteria, which are based on a single preference order on strategy profiles, are *cooperative*, in that they capture social aspects of rational behavior. We define and study, for the first time, an *adversarial* external criterion for rational behavior. Put simply, we assume that each player attempts to minimize the other player's payoff as long as, by doing so, she does not decrease her own payoff. This yields two different preference orders on strategy profiles, one for each player, and gives rise to a new notion of equilibrium.

Adversarial External Choice. According to our notion of rationality, among two strategy profiles (σ, π) and (σ', π'), player 1 prefers (σ, π), denoted $(\sigma, \pi) \succeq_1 (\sigma', \pi')$, if either $v_1^{\sigma,\pi} > v_1^{\sigma',\pi'}$, or both $v_1^{\sigma,\pi} = v_1^{\sigma',\pi'}$ and $v_2^{\sigma,\pi} \leq v_2^{\sigma',\pi'}$. In other words, the *preference order* $\succeq_1$ of player 1 is lexicographic: the primary goal of player 1 is to maximize her own payoff; the secondary goal is to minimize the opponent's payoff. The preference order $\succeq_2$ of player 2 is defined symmetrically. It should be noted that, defined in this way, adversarial external choice cannot be internalized uniformly over all games by changing the payoff functions of the two players: if $v_1^{\sigma,\pi} = v_1^{\sigma',\pi'}$ and $v_2^{\sigma,\pi} \leq v_2^{\sigma',\pi'}$, then uniform internalization would require to increase $v_1^{\sigma,\pi}$ by an arbitrarily small $\varepsilon > 0$.

Secure Equilibria. The two orders $\succeq_1$ and $\succeq_2$ on strategy profiles, which express the preferences of the two players, induce the following refinement of the Nash equilibrium notion: (σ, π) is a *secure equilibrium* if (1) $(v_1^{\sigma,\pi}, v_2^{\sigma,\pi}) \succeq_1 (v_1^{\sigma',\pi}, v_2^{\sigma',\pi})$ for all player 1 strategies σ', and (2) $(v_1^{\sigma,\pi}, v_2^{\sigma,\pi}) \succeq_2 (v_1^{\sigma,\pi'}, v_2^{\sigma,\pi'})$ for all player 2 strategies π'. Note that every secure equilibrium is a Nash equilibrium, but a Nash equilibrium need not be secure. The name "secure" equilibrium derives from the following equivalent characterization. We say that a strategy profile (σ, π) is *secure* if any rational deviation of player 2 —i.e., a deviation that does not decrease her payoff— will not decrease the payoff of player 1, and symmetrically, any rational deviation of player 1 will not decrease the payoff of player 2. Formally, (σ, π) is secure if for all player 2 strategies π', if $v_2^{\sigma,\pi'} \geq v_2^{\sigma,\pi}$ then $v_1^{\sigma,\pi'} \geq v_1^{\sigma,\pi}$, and for all player 1 strategies σ', if $v_1^{\sigma',\pi} \geq v_1^{\sigma,\pi}$ then $v_2^{\sigma',\pi} \geq v_2^{\sigma,\pi}$. The secure profile (σ, π) can thus be interpreted as a contract between the two players which enforces cooperation: any unilateral selfish deviation by one player cannot put the other player at a disadvantage if she follows the contract. It is not difficult to show (see Section 2) that a strategy profile is a secure equilibrium iff it is both a secure profile and a Nash equilibrium. Thus, the secure equilibria are those Nash equilibria which represent enforceable contracts between the two players.

Motivation: Verification of Component-Based Systems. The motivation for our definitions comes from verification. There, one would like to prove that a component of a system (player 1) can satisfy a specification no matter how the environment (player 2) behaves [3]. Classically, this is modeled as a strictly competitive (zero-sum) game, where the environment's objective is the complement of the component's objective. However, the zero-sum model is often naive, as the environment itself typically consists of components, each with its own specification (i.e., objective). Moreover, the individual component specifications are usually not complementary; a common example is that each component must maintain a local invariant. So a more appropriate approach is to prove that player 1 can meet her objective no matter how player 2 behaves *as long as* player 2 does not sabotage her own objective. In other words, classical correctness proofs of a component assume absolute worst-case behavior of the environment, while it would suffice to assume only *relative* worst-case behavior of the environment —namely, relative to the assumption that the environment itself is correct (i.e., meets its specification). Such relative worst-case reasoning called assume-guarantee reasoning [1,2,13] so far has not been studied in the natural setting offered by game theory.

Existence and Uniqueness of Maximal Secure Equilibria. We will see that in general games, such as matrix games, there may be multiple secure equilibrium payoff profiles, even several incomparable maximal ones. However, the games that occur in verification have a special form. They are played on directed graphs whose nodes represent system states, and whose edges represent system transitions. The nodes partitioned into two sets: in player 1 nodes, the first player chooses an outgoing edge, and in player 2 nodes, the second player chooses an outgoing edge. By repeating these choices ad infinitum, an infinite path through the graph is formed, which represents a system trace. The objective φ_i of each player i is a set of infinite paths; for example, an invariant (or "safety") objective is the set of infinite paths that do not visit unsafe states. Each player i attempts to satisfy her objective φ_i by choosing a strategy that ensures that the outcome of the game lies in the set φ_i. The objective φ_i is typically an ω-regular set (specified, e.g., in temporal logic), or more generally, a Borel set [8] in the Cantor topology on infinite paths. We call these games 2-player non-zero-sum *graph games with Borel objectives*. Our main result shows that for these games, which may have multiple maximal Nash equilibria, there always exists a unique maximal secure equilibrium payoff profile. In other words, in graph games with Borel objectives there is a compelling notion of rational behavior for each player, which is (1) a classical Nash equilibrium, (2) an enforceable contract ("secure"), and (3) a guarantee of maximal payoff for each player among all behaviors that achieve (1) and (2).

Examples. Consider the game graph shown in Fig. 1. Player 1 chooses the successor node at square nodes and her objective is to reach the target s_4, a reachability (co-safety) objective. Player 2 chooses the successor node at diamond nodes and her objective is to reach s_3 or s_4, also a reachability objective. There are two player 1 strategies: σ_1 chooses the move $s_0 \to s_1$, and σ_2 chooses $s_0 \to s_2$.

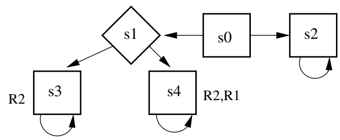

Fig. 1. A graph game with reachability objectives

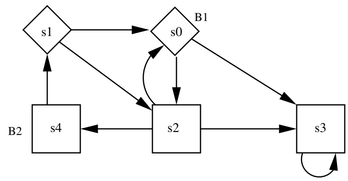

Fig. 2. A graph game with Büchi objectives

There are also two player 2 strategies: π_1 chooses $s_1 \to s_3$, and π_2 chooses $s_1 \to s_4$. The strategy profile (σ_1, π_1) leads the game into s_3 and therefore gives the payoff profile (0,1), meaning player 1 loses and player 2 wins (i.e., only player 2 reaches her target). The strategy profiles (σ_1, π_2), (σ_2, π_1), and (σ_2, π_2) give the payoffs (1,1), (0,0), and (0,0), respectively. All four strategy profiles are Nash equilibria; for example, in (σ_1, π_1) player 1 does not have an incentive to switch to strategy σ_2 (which would still give her payoff 0), and neither does player 2 have an incentive to switch to π_2 (she is already getting payoff 1). However, the strategy profile (σ_1, π_1) is not a secure equilibrium, because player 2 can lower player 1's payoff (from 1 to 0) without changing her own payoff by switching to strategy σ_2. Similarly, the strategy profile (σ_1, π_2) is not secure, because player 1 can lower player 2's payoff without changing her own payoff by switching to σ_1. So if both players, in addition to maximizing their own payoff, also attempt to minimize the opponents payoff, then the resulting payoff profile is unique, namely, (0,0). In other words, in this game, the only rational behavior for both players is to deny each other's objectives.

This is not always the case: sometimes it is beneficial for both players to cooperate to achieve their own objectives, with the result that both players win. Consider the game graph shown in Fig. 2. Both players have Büchi objectives: player 1 (square) wants to visit s_0 infinitely often, and player 2 (diamond) wants to visit s_4 infinitely often. If player 2 always chooses $s_1 \to s_0$ and player 1 always chooses $s_2 \to s_4$, then both players win. This Nash equilibrium is also secure: if player 1 deviates by choosing $s_2 \to s_0$, then player 2 can "retaliate" by choosing $s_0 \to s_3$; similarly, if player 2 deviates by choosing $s_1 \to s_2$, then player 2 can retaliate by $s_2 \to s_3$. It follows that for purely selfish motives (and not some social reason), both players have an incentive to cooperate to achieve the maximal secure equilibrium payoff (1,1).

Outline and Results. In Section 2, we define the notion of secure equilibrium and give several interpretations through alternative definitions. In Section 3 we prove the existence and uniqueness of maximal secure equilibria in graph games with Borel objectives. The proof is based on the following classification of strategies. A player 1 strategy is called *strongly winning* if it ensures that player 1 wins and player 2 loses (i.e., the outcome of the game satisfies $\varphi_1 \wedge \neg\varphi_2$). A player 1 strategy is *retaliating* if it ensures that player 1 wins if player 2 wins (i.e., the outcome satisfies $\varphi_2 \to \varphi_1$). In other words, a retaliating strategy for player 1 ensures that if player 2 causes player 1 to lose, then player 2 will lose too. If both players follow retaliating strategies (σ, π), they may both win —in this case, we say that (σ, π) is a *winning pair of retaliating strategies*— or they may both lose. We show that at every node of a graph game with Borel objectives, either one of the two players has a strongly winning strategy, or there is a pair of retaliating strategies. Based on this insight, we give an algorithm for computing the secure equilibria in graph games in the case that both players' objectives are ω-regular. In Section 4, we analyze the memory requirements of strongly winning and retaliating strategies in graph games with ω-regular objectives. Our results (in Table 1 and 2) consider safety, reachability, Büchi, co-Büchi, and general parity objectives. We show that strongly winning and retaliating strategies often require memory, even in the simple case that a player pursues a reachability objective. In Section 5, we generalize the notion of secure equilibria from 2-player to n-player games. We show that there can be multiple maximal secure equilibria in 3-player graph games with reachability objectives.

2 Definitions

In a secure game the objective of player 1 is to maximize her own payoff and then minimize the payoff of player 2. Similarly, player 2 maximizes her own payoff and then minimizes the payoff of player 1. We want to determine the best payoff that each player can ensure when both players play according to these preferences. We formalize this as follows. A *strategy profile* (σ, π) is a pair of strategies, where σ is a player 1 strategy and π is a player 2 strategy. The strategy profile (σ, π) gives rise to a *payoff profile* $(v_1^{\sigma,\pi}, v_2^{\sigma,\pi})$, where $v_1^{\sigma,\pi}$ is the payoff of player 1 if the two players follow the strategies σ and π respectively, and $v_2^{\sigma,\pi}$ is the corresponding payoff of player 2. We define the player 1 preference order $\preceq_1$ and the player 2 preference order $\preceq_2$ on payoff profiles lexicographically:

$$(v_1, v_2) \prec_1 (v_1', v_2') \text{ iff } (v_1 < v_1') \vee (v_1 = v_1' \wedge v_2 > v_2'),$$

that is, player 1 prefers a payoff profile which gives her greater payoff, and if two payoff profiles match in the first component, then she prefers the payoff profile in which the player 2's payoff is minimized; symmetrically,

$$(v_1, v_2) \prec_2 (v_1', v_2') \text{ iff } (v_2 < v_2') \vee (v_2 = v_2' \wedge v_1 > v_1').$$

Given two payoff profiles (v_1, v_2) and (v_1', v_2'), we write $(v_1, v_2) = (v_1', v_2')$ iff $v_1 = v_1'$ and $v_2 = v_2'$, and $(v_1, v_2) \preceq_1 (v_1, v_2')$ iff either $(v_1, v_2) \prec_1 (v_1', v_2')$ or $(v_1, v_2) = (v_1', v_2')$. We define $\preceq_2$ analogously.

Definition 1 (Secure strategy profiles). *A strategy profile* (σ, π) *is* secure *if the following two conditions hold:*

$$\forall \pi'.\ (v_1^{\sigma,\pi'} < v_1^{\sigma,\pi}) \rightarrow (v_2^{\sigma,\pi'} < v_2^{\sigma,\pi})$$

$$\forall \sigma'.\ (v_2^{\sigma',\pi} < v_2^{\sigma,\pi}) \rightarrow (v_1^{\sigma',\pi} < v_1^{\sigma,\pi})$$ ∎

A secure strategy for player 1 ensures that if player 2 tries to decrease player 1's payoff, then player 2's payoff decreases as well, and vice versa.

Definition 2 (Secure equilibria). *A strategy profile* (σ, π) *is a* secure equilibrium *if the strategy profile is a Nash equilibrium and it is secure.* ∎

Lemma 1 (Equivalent characterization). *The strategy profile* (σ, π) *is a secure equilibrium iff the following two conditions hold:*

$$\forall \pi'.\ (v_1^{\sigma,\pi'}, v_2^{\sigma,\pi'}) \preceq_2 (v_1^{\sigma,\pi}, v_2^{\sigma,\pi})$$

$$\forall \sigma'.\ (v_1^{\sigma',\pi}, v_2^{\sigma',\pi}) \preceq_1 (v_1^{\sigma,\pi}, v_2^{\sigma,\pi})$$

Proof. Given (σ, π) is a Nash equilibrium strategy profile we have for all π', $v_2^{\sigma,\pi'} \leq v_2^{\sigma,\pi}$. Since the strategy profile is also a secure strategy profile for all strategy π' we have $(v_1^{\sigma,\pi'} < v_1^{\sigma,\pi}) \rightarrow (v_2^{\sigma,\pi'} < v_2^{\sigma,\pi})$. It follows from above that for any arbitrary π' the following condition hold:

$$(v_2^{\sigma,\pi'} = v_2^{\sigma,\pi}\ \wedge\ v_1^{\sigma,\pi} \leq v_1^{\sigma,\pi'}) \vee (v_2^{\sigma,\pi'} < v_2^{\sigma,\pi}).$$

Hence for all π' we have $(v_1^{\sigma,\pi'}, v_2^{\sigma,\pi'}) \preceq_2 (v_1^{\sigma,\pi}, v_2^{\sigma,\pi})$. The argument for the other case is symmetric.

Hence neither player 1 nor player 2 has any incentive to switch from the strategy profile (σ, π) to increase the payoff profile according to their respective payoff profile ordering. ∎

Example 1 (Matrix games). A secure equilibrium need not exist in a matrix game. We give an example of a matrix game where no Nash equilibrium is secure. Consider the game M_1 below, where the row player can choose row 1 or row 2 (denoted r_1 and r_2, respectively), and the column player chooses between the two columns (denoted c_1 and c_2). The first component of the payoff is the row player payoff, and the second component is the column player payoff.

$$M_1 \;=\; \begin{bmatrix} (3,3) & (1,3) \\ (3,1) & (2,2) \end{bmatrix}$$

In this game the strategy profile (r_1, c_1) is the only Nash equilibrium. But (r_1, c_1) is not a secure strategy profile, because if the row player plays r_1, then the column player playing c_2 can still get payoff 3 and decrease the row player's payoff to 1. In the game M_2 there are two Nash equilibria, namely, (r_1, c_2) and (r_2, c_1), and the strategy profile (r_2, c_1) is a secure strategy profile as well. Hence the strategy profile (r_2, c_1) is a secure equilibrium. However the strategy profile (r_1, c_2) is not secure.

$$M_2 = \begin{bmatrix} (0,0) & (1,0) \\ (\frac{1}{2},\frac{1}{2}) & (\frac{1}{2},\frac{1}{2}) \end{bmatrix}$$

Multiple secure equilibria can exist, as in the case, for example, in a matrix game where all entries of the matrix are the same. We now present an example of a matrix game with multiple secure equilibria profile. Consider the following matrix game M_3. The strategy profile (r_1, c_1) and (r_2, c_2) are both secure equilibria. The former has a payoff profile $(2,1)$ and the later has a payoff profile $(1,2)$. Hence there can be multiple secure equilibria payoff profiles and in case there are multiple secure equilibria payoff profiles the maximal payoff profile is not always unique.

$$M_3 = \begin{bmatrix} (2,1) & (0,0) \\ (0,0) & (1,2) \end{bmatrix} \qquad \blacksquare$$

3 2-Player Non-zero-sum Games on Graphs

We consider 2-player infinite path-forming games played on graphs. A *game graph* $G = ((V, E), (V_1, V_2))$ consists of a directed graph (V, E), where V is the set of states (vertices) and E is the set of edges, and a partition (V_1, V_2) of the states. For technical convenience we assume that every state has at least one outgoing edge. The two players, player 1 and player 2, keep moving a token along the edges of the game graph: player 1 moves the token from states in V_1, and player 2 moves the token from states in V_2. A *play* is an infinite path $\Omega = \langle s_0, s_1, s_2, \ldots \rangle$ through the game graph, that is, $(s_k, s_{k+1}) \in E$ for all $k \geq 0$. A strategy for player 1, given a prefix of a play (i.e., a finite sequence of states), specifies a next state to extend the play. Formally, a *strategy* for player 1 is a function σ: $V^* \cdot V_1 \to V$ such that for all $x \in V^*$ and $s \in V_1$, we have $(s, \sigma(x \cdot s)) \in E$. A strategy π for player 2 is defined symmetrically. We write Σ and Π to denote the sets of strategies for player 1 and player 2, respectively. A strategy is memoryless if it is independent of the history of play. Formally, a strategy τ of player i, where $i \in \{1, 2\}$, is *memoryless* if $\tau(x \cdot s) = \tau(x' \cdot s)$ for all $x, x' \in V^*$ and all $s \in V_i$; hence a memoryless strategy of player i can be represented as a function τ: $V_i \to V$. A play $\Omega = \langle s_0, s_1, s_2, \ldots \rangle$ is *consistent* with a strategy τ of player i if for all $k \geq 0$, if $s_k \in V_i$, then $s_{k+1} = \tau(s_0, s_1, \ldots, s_k)$. Given a state $s \in V$, a strategy σ of player 1, and a strategy π of player 2, there is a unique play $\Omega_{\sigma,\pi}(s)$, the *outcome* of the game, which starts from s and is consistent with both σ and π.

Objectives of the players are specified generally as sets $\varphi \subseteq V^\omega$ of infinite paths. We write $\Omega \models \varphi$ instead of $\Omega \in \varphi$ for infinite paths Ω and objectives φ. We use boolean operators such as $\vee$, $\wedge$, and $\neg$ on objectives to denote set union, intersection, and complement. A *Borel objective* is a Borel set $\varphi \subseteq V^\omega$ in the Cantor topology on V^ω. The following celebrated result of Martin establishes that all games with Borel objectives are determined.

Theorem 1 (Borel determinacy [11]). *For every 2-player graph game G, every state s, and every Borel objective φ, either (1) there is a strategy σ of* *player 1 such that for all strategies* π' *of player 2, we have* $\Omega_{\sigma,\pi'}(s) \models \varphi$*, or (2) there is a strategy* π *of player 2 such that for all strategies* σ' *of player 1, we have* $\Omega_{\sigma',\pi}(s) \models \neg\varphi$.

In verification, objectives are usually *ω-regular* sets. The ω-regular sets occur in the low levels of the Borel hierarchy (in $\Sigma_3 \cap \Pi_3$) and form a robust and expressive class for determining the payoffs of commonly used system specifications [10,16].

We consider non-zero-sum games on graphs. For our purposes, a *graph game* $(G, s, \varphi_1, \varphi_2)$ consists of a game graph G, say with state set V, together with a start state $s \in V$ and two Borel objectives $\varphi_1, \varphi_2 \subseteq V^\omega$. The game starts at state s, player 1 pursues the objective φ_1, and player 2 pursues the objective φ_2 (in general, φ_2 is not the complement of φ_1). Player $i \in \{1,2\}$ gets payoff 1 if the outcome of the game is a member of φ_i, and she gets payoff 0 otherwise. In the following, we fix the game graph G and the objectives φ_1 and φ_2, but we vary the start state s of the game. Thus we parameterize the payoffs by s: given strategies σ and π for the two players, we write $v_i^{\sigma,\pi}(s) = 1$ if $\Omega_{\sigma,\pi}(s) \models \varphi_i$, and $v_i^{\sigma,\pi}(s) = 0$ otherwise, for $i \in \{1,2\}$. Similarly, we sometimes refer to Nash equilibria and secure strategy profiles of the graph game $(G, s, \varphi_1, \varphi_2)$ as equilibria and secure profiles *at* the state s.

3.1 Unique Maximal Secure Equilibria

Consider a game graph G with state set V, and Borel objectives φ_1 and φ_2 for the two players.

Definition 3 (Maximal secure equilibria). *For* $v, w \in \{0,1\}$*, we write* $S_{v,w} \subseteq V$ *to denote the set of states* s *such that a secure equilibrium with the payoff profile* (v,w) *exists in the game* $(G, s, \varphi_1, \varphi_2)$*, that is,* $s \in S_{v,w}$ *iff there is a secure equilibrium* (σ,π) *at* s *such that* $(v_1^{\sigma,\pi}(s), v_2^{\sigma,\pi}(s)) = (v,w)$*. Similarly,* $MS_{v,w} \subseteq S_{v,w}$ *denotes the set of states* s *such that the payoff profile* (v,w) *is a* maximal secure equilibrium payoff profile *at* s*, that is,* $s \in MS_{v,w}$ *iff (1)* $s \in S_{v,w}$ *and (2) for all* $v', w' \in \{0,1\}$*, if* $s \in S_{v',w'}$*, then* $(v',w') \preceq_1 (v,w)$ *and* $(v',w') \preceq_2 (v,w)$*.* ∎

We now define the notions of strongly winning and retaliating strategies, which capture the essence of secure equilibria. A strategy for player 1 is strongly winning if it ensures that the objective of player 1 is satisfied and the objective of player 2 is not. A retaliating strategy for player 1 ensures that for every strategy of player 2, if the objective of player 2 is satisfied, then the objective of player 1 is satisfied as well. We will show that every secure equilibrium either contains a strongly winning strategy for one of the players, or it consists of a pair of retaliating strategies.

Definition 4 (Strongly winning strategies). *A strategy* σ *is* strongly winning *for player 1 from a state* s *if she can ensure the payoff profile* $(1,0)$ *in the game* $(G, s, \varphi_1, \varphi_2)$ *by playing the strategy* σ*. Formally,* σ *is strongly winning* *for player 1 if for all player 2 strategies π, we have $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \neg\varphi_2)$. The strongly winning strategies for player 2 are defined symmetrically.* ∎

Definition 5 (Retaliating strategies). *A strategy σ is a* retaliating *strategy for player 1 from a state s if for all player 2 strategies π, we have $\Omega_{\sigma,\pi}(s) \models (\varphi_2 \rightarrow \varphi_1)$. Similarly, a strategy π is a retaliating strategy for player 2 from s if for all player 1 strategies σ, we have $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \rightarrow \varphi_2)$. We write $Re_1(s)$ and $Re_2(s)$ to denote the sets of retaliating strategies for player 1 and player 2 from s. A strategy profile (σ, π) is a* retaliation strategy profile *at a state s if both σ and π are retaliating strategies from s.* ∎

Example 2 (Büchi-Büchi game). Recall the game shown in Fig. 2. Consider the memoryless strategies of player 2 at state s_0. If player 2 chooses $s_0 \rightarrow s_3$, then player 2 does not satisfy her Büchi objective. If player 2 chooses $s_0 \rightarrow s_2$, then at state s_2 player 1 chooses $s_2 \rightarrow s_0$, and hence player 1's objective is satisfied, but player 2's objective is not satisfied. Thus, no memoryless strategy for player 2 can be a winning retaliating strategy at s_0.

Now consider the strategy π_g for player 2 which chooses $s_0 \rightarrow s_2$ if between the last two consecutive visits to s_0 the state s_4 was visited, and otherwise it chooses $s_0 \rightarrow s_3$. Given this strategy, for every strategy of player 1 that satisfies player 1's objective, player 2's objective is also satisfied. Let σ_g be the player 1 strategy that chooses $s_2 \rightarrow s_4$ if between the last two consecutive visits to s_2 the state s_0 was visited, and otherwise chooses $s_2 \rightarrow s_3$. The strategy profile (σ_g, π_g) consists of a pair of winning retaliating strategies, as it satisfies the Büchi objectives of both players. If instead, player 2 always chooses $s_0 \rightarrow s_3$, and player 1 always chooses $s_2 \rightarrow s_3$, we obtain a memoryless retaliation strategy profile, which is not winning for either player: it is a Nash equilibrium at state s_0 with the payoff profile $(0, 0)$. Finally, suppose that at s_0 player 2 always chooses s_2, and at s_2 player 1 always chooses s_0. This strategy profile is again a Nash equilibrium, with the payoff profile $(0, 1)$ at s_0, but not a retaliation strategy profile. This shows that at state s_0 the Nash equilibrium payoff profiles $(0, 1)$, $(0, 0)$, and $(1, 1)$ are possible, but only $(0, 0)$ and $(1, 1)$ are secure. ∎

Definition 6 (Winning sets). *We define the following state sets in terms of strongly winning and retaliating strategies.*

- *The sets of states where player 1 or player 2 have a strongly winning strategy, denoted by W_{10} and W_{01}, respectively:*

$$W_{10} = \{ s \in V \,:\, \exists \sigma \in \Sigma.\ \forall \pi \in \Pi.\ \ \Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \neg\varphi_2) \}$$

$$W_{01} = \{ s \in V \,:\, \exists \pi \in \Pi.\ \forall \sigma \in \Sigma.\ \ \Omega_{\sigma,\pi}(s) \models (\varphi_2 \wedge \neg\varphi_1) \}$$

- *The set of states where both players have retaliating strategies and there exists a retaliation strategy profile whose strategies satisfy the objectives of both players:*

$$W_{11} = \{ s \in V \,:\, \exists \sigma \in Re_1(s).\ \exists \pi \in Re_2(s).\ \ \Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \varphi_2) \}$$

– *The set of states where both players have retaliating strategies and for every retaliation strategy profile, neither the objective of player 1 nor the objective of player 2 is satisfied:*

$$W_{00} = \{\, s \in V \;:\; Re_1(s) \neq \emptyset \;\; and \;\; Re_2(s) \neq \emptyset \; and \\ \forall \sigma \in Re_1(s).\, \forall \pi \in Re_2(s).\;\; \Omega_{\sigma,\pi}(s) \models (\neg\varphi_1 \wedge \neg\varphi_2) \,\}$$ ∎

We show that the four sets W_{10}, W_{01}, W_{11}, and W_{00} form a partition of the state space. This result fully characterizes each state of a 2-player non-zero-sum graph game with Borel objectives, just like the determinacy result (Theorem 1) fully characterizes the zero-sum case. In the zero-sum case, where $\varphi_2 = \neg\varphi_1$, the sets W_{10} and W_{01} specify the winning states for players 1 and 2, respectively, $W_{11} = \emptyset$ by definition, and $W_{00} = \emptyset$ by determinacy. We also show that for all $v, w \in \{0,1\}$, we have $MS_{v,w} = W_{v,w}$. It follows that for 2-player graph games (1) secure equilibria always exist, and moreover, (2) there is always a unique maximal secure equilibrium payoff profile. (Example 2 showed that there can be multiple secure equilibria with different payoff profiles). The proof proceeds in several steps.

Lemma 2. $W_{10} = \{\, s \in V \;:\; Re_2(s) = \emptyset \,\}$ *and* $W_{01} = \{\, s \in V \;:\; Re_1(s) = \emptyset \,\}$.

Proof. We show the inclusion of one set in the other for both the direction:

1. $W_{10} \subseteq \{\, s : Re_2(s) = \emptyset \,\}$ as a strongly winning strategy σ of player 1 to satisfy $(\varphi_1 \wedge \neg\varphi_2)$ against any strategy π of player 2 is a witness to exhibit that there is no retaliation strategy for player 2.
2. It follows from Borel determinacy (Theorem 1) that from every state s in $V \setminus W_{10}$ there is a strategy π for player 2 to satisfy $(\neg\varphi_1 \vee \varphi_2)$ against any strategy of player 1. The strategy π is a retaliation strategy for player 2. Hence we have $V \setminus W_{10} \subseteq \{\, s : Re_2(s) \neq \emptyset \,\}$.

The claim is a consequence of the above facts. ∎

Lemma 3. *Consider the following sets:*

$$T_1 = \{\, s \in V \;:\; \forall \sigma \in Re_1(s).\, \forall \pi \in Re_2(s).\;\; \Omega_{\sigma,\pi}(s) \models (\neg\varphi_1 \wedge \neg\varphi_2) \,\}$$

$$T_2 = \{\, s \in V \;:\; \forall \sigma \in Re_1(s).\, \forall \pi \in Re_2(s).\;\; \Omega_{\sigma,\pi}(s) \models (\neg\varphi_1 \vee \neg\varphi_2) \,\}$$

Then $T_1 = T_2$.

Proof. The inclusion $T_1 \subseteq T_2$ follows from the fact that $(\neg\varphi_1 \wedge \neg\varphi_2) \rightarrow (\neg\varphi_1 \vee \neg\varphi_2)$. We show that $T_2 \subseteq T_1$. By the definition of retaliating strategies, if σ is a retaliating strategy of player 1, then for all strategies π of player 2, we have $\Omega_{\sigma,\pi}(s) \models (\varphi_2 \rightarrow \varphi_1)$, and thus $\Omega_{\sigma,\pi}(s) \models (\neg\varphi_1 \rightarrow \neg\varphi_2)$. Symmetrically, if π is a retaliating strategy of player 2, then for all strategies σ of player 1, we have $\Omega_{\sigma,\pi}(s) \models (\neg\varphi_2 \rightarrow \neg\varphi_1)$. The claim follows. ∎

It follows from Lemma 2 and Lemma 3 that $W_{00} = V \setminus (W_{01} \cup W_{10} \cup W_{11})$. It also follows from Lemma 2 that the sets W_{01}, W_{10}, and W_{11} are disjoint. This gives the following result.

Theorem 2 (State space partition). *For all 2-player graph games with Borel objectives, the four sets W_{10}, W_{01}, W_{11}, and W_{00} form a partition of the state set.*

Lemma 4. *Consider the sets S_{ij} for $i, j \in \{ 0, 1 \}$ as defined in Definition 3. The following equalities hold:*

$$S_{00} \cap S_{01} = \emptyset; \quad S_{00} \cap S_{10} = \emptyset;$$

$$S_{01} \cap S_{10} = \emptyset; \quad S_{11} \cap S_{01} = \emptyset; \quad S_{11} \cap S_{10} = \emptyset.$$

Proof. Consider a state $s \in S_{10}$ and a secure equilibrium (σ, π) at s. Since the strategy profile is secure and player 2 gets the least possible payoff, it follows that for all player 1 strategies π', the payoff for player 1 cannot decrease. Hence for all player 2 strategies π', we have $\Omega_{\sigma,\pi'}(s) \models \varphi_1$. So there is no Nash equilibrium at state s which assigns payoff 0 to player 1. Hence we have $S_{10} \cap S_{01} = \emptyset$ and $S_{10} \cap S_{00} = \emptyset$. By symmetry, $S_{01} \cap S_{00} = \emptyset$.

Consider a state $s \in S_{11}$ and a secure equilibrium (σ, π) at s. Since the strategy profile is secure, it ensures that for all player 2 strategies π', if $\Omega_{\sigma,\pi'}(s) \models \neg\varphi_1$, then $\Omega_{\sigma,\pi'} \models \neg\varphi_2$. Hence $s \notin S_{01}$. Thus we have $S_{11} \cap S_{01} = \emptyset$, and by symmetry $S_{11} \cap S_{10} = \emptyset$. ∎

Lemma 5. *The following equalities hold:*

$$MS_{00} \cap MS_{01} = \emptyset; \quad MS_{00} \cap MS_{10} = \emptyset;$$

$$MS_{01} \cap MS_{10} = \emptyset; \quad MS_{11} \cap MS_{00} = \emptyset.$$

Proof. The first three equalities follow from Lemma 4. The last equality follows from the facts that $(0, 0) \preceq_1 (1, 1)$ and $(0, 0) \preceq_2 (1, 1)$. So if $s \in MS_{11}$, then $(0, 0)$ cannot be a maximal secure payoff profile at s. ∎

Lemma 6. $W_{11} = MS_{11}$*;* $W_{10} = MS_{10}$*;* $W_{01} = MS_{01}$*.*

Proof. Consider a state $s \in MS_{10}$ and a secure equilibrium (σ, π) at s. Since player 2 gets the least possible payoff and (σ, π) is a secure strategy profile, it follows that for all strategies π' of player 2, we have $\Omega_{\sigma,\pi'}(s) \models \varphi_1$. Since (σ, π) is a Nash equilibrium, for all strategies π' of player 2, we have $\Omega_{\sigma,\pi'}(s) \models \neg\varphi_2$. Thus we have $MS_{10} \subseteq W_{10}$. Now consider a state $s \in W_{10}$ and let σ be a strongly winning strategy of player 1 at s, that is, for all strategies π of player 2, we have $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \neg\varphi_2)$. For all strategies π of player 2, the strategy profile (σ, π) is a secure equilibrium. Hence, $s \in S_{10}$. Since $(1, 0)$ is the greatest payoff profile in the preference ordering of the payoff profiles for player 1, we have $s \in MS_{10}$. Therefore $W_{10} = MS_{10}$. Symmetrically, $W_{01} = MS_{01}$.

Consider a state $s \in MS_{11}$ and let (σ, π) be a secure equilibrium at s. We prove that $\sigma \in Re_1(s)$ and $\pi \in Re_2(s)$. Since (σ, π) is a secure strategy profile, for all strategies π' of player 2, if $\Omega_{\sigma,\pi'}(s) \models \neg\varphi_1$, then $\Omega_{\sigma,\pi'}(s) \models \neg\varphi_2$. In other words, for all strategies π' of player 2, we have $\Omega_{\sigma,\pi'}(s) \models (\varphi_2 \rightarrow \varphi_1)$. Hence $\sigma \in Re_1(s)$. Symmetrically, $\pi \in Re_2(s)$. Thus $MS_{11} \subseteq W_{11}$. Consider a state

$s \in W_{11}$ and let $\sigma \in Re_1(s)$ and $\pi \in Re_2(s)$ such that $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \varphi_2)$. A retaliation strategy profile is, by definition, a secure strategy profile. Since the strategy profile (σ, π) assigns the greatest possible payoff to each player, it is a Nash equilibrium. Therefore $W_{11} \subseteq S_{11} \subseteq MS_{11}$. ∎

Lemma 7. $W_{00} = MS_{00}$.

Proof. It follows from Lemma 4 and Lemma 5 that $MS_{00} = S_{00} \setminus S_{11} = S_{00} \setminus MS_{11}$. We will use this fact to prove that $W_{00} = MS_{00}$.

- Consider a state $s \in MS_{00}$. Then we have $s \notin (MS_{11} \cup MS_{10} \cup MS_{11}) \Rightarrow s \notin (W_{11} \cup W_{10} \cup W_{01})$. By Lemma 2 we have $W_{00}, W_{11}, W_{10}, W_{01}$ is a partition and hence we have $s \in W_{00}$. It follows that $MS_{00} \subseteq W_{00}$.
- Consider a state $s \in W_{00}$. We claim that there is a strategy σ for player 1 such that for all strategy π' we have $\Omega_{\sigma,\pi'}(s) \models \neg\varphi_2$. Assume by the way of contradiction this is not the case. By Borel determinacy then we have there is a strategy π'' for player 2 such that for all σ' we have $\Omega_{\sigma',\pi''}(s) \models \varphi_2$. It follows that either π'' is a strongly winning strategy for player 2 or a retaliation strategy such that player 2 gets payoff 1. Hence $s \notin W_{00}$, which is a contradiction. Hence there is a strategy σ such that for all π' we have $\Omega_{\sigma,\pi'}(s) \models \neg\varphi_2$. Similarly, there is a strategy π such that for all σ' we have $\Omega_{\sigma',\pi}(s) \models \neg\varphi_1$. We claim that (σ, π) is a secure equilibrium strategy profile. By property of σ for any π', $\Omega_{\sigma,\pi'}(s) \models \neg\varphi_2$. Similar argument hold for π as well. Hence we have (σ, π) is a Nash equilibrium strategy profile. For the strategy profile (σ, π) we have the payoff profile is $(0, 0)$ and it assigns the least possible payoff to each player. Hence it is a secure strategy profile. Hence $s \in S_{00}$. Also $s \in W_{00} \Rightarrow s \notin W_{11} = MS_{11}$. Hence $s \in S_{00} \setminus MS_{11}$. This gives us $W_{00} \subseteq MS_{00}$. ∎

Theorem 2 together with Lemmas 6 and 7 yields the following result.

Theorem 3 (Unique maximal secure equilibrium). *At every state of a 2-player graph game with Borel objectives, there exists a unique maximal secure equilibrium payoff profile.*

3.2 Algorithmic Characterization

We now give an alternative characterization of the sets W_{00}, W_{01}, W_{10}, and W_{11}. The new characterization is useful to derive computational complexity results for computing the four sets when player 1 and player 2 have ω-regular objectives. The characterization itself, however, is general and applies to all objectives specified as Borel sets.

Definition 7 (Cooperative strategy profiles). *Given a game graph G with state set V, and an objective $\psi \subseteq V^\omega$, we define the following sets:*

$$\langle\langle 1 \rangle\rangle_G \, \psi = \{s \in V \, : \, \exists \sigma \in \Sigma. \; \forall \pi \in \Pi. \;\; \Omega_{\sigma,\pi}(s) \models \psi\}$$

$$\langle\langle 2\rangle\rangle_G\,\psi = \{s \in V \,:\, \exists\pi \in \Pi.\ \forall\sigma \in \Sigma.\ \ \Omega_{\sigma,\pi}(s) \models \psi\}$$

$$\langle\langle 1,2\rangle\rangle_G\,\psi = \{s \in V \,:\, \exists\sigma \in \Sigma.\ \exists\pi \in \Pi.\ \ \Omega_{\sigma,\pi}(s) \models \psi\}$$

We omit the subscript G if it is clear from the context. Let s be a state in $\langle\langle 1,2\rangle\rangle\psi$ *and let* (σ,π) *be a strategy profile such that* $\Omega_{\sigma,\pi}(s) \models \psi$. *We refer to* (σ,π) *as a* cooperative *strategy profile at s, and informally say that the two players are cooperating to satisfy* ψ. ∎

It follows from the definitions that $W_{10} = \langle\langle 1\rangle\rangle(\varphi_1 \wedge \neg\varphi_2)$ and $W_{01} = \langle\langle 2\rangle\rangle(\varphi_2 \wedge \neg\varphi_1)$. Define $A = V \setminus (W_{10} \cup W_{01})$, the set of "ambiguous" states from which neither player has a strongly winning strategy. Let $W_i = \langle\!\langle i\rangle\!\rangle\varphi_i$, for $i \in \{\,1,2\,\}$, the winning sets of the two players, and let $U_1 = W_1 \setminus W_{10}$ and $U_2 = W_2 \setminus W_{01}$, the sets of "weakly winning" states for players 1 and 2, respectively. Define $U = U_1 \cup U_2$. Note that $U \subseteq A$.

Lemma 8. $U \subseteq W_{11}$.

Proof. Let $s \in U_1$. By the definition of U_1, player 1 has a strategy σ from the state s to satisfy the objective φ_1, which is obviously a retaliating strategy, because φ_1 implies $\varphi_2 \rightarrow \varphi_1$. Again by the definition of U_1, we have $s \notin W_{10}$. Hence, by the determinacy of zero-sum games (Theorem 1), player 2 has a strategy π to satisfy the objective $\neg(\varphi_1 \wedge \neg\varphi_2)$, which is a retaliating strategy, because $\neg(\varphi_1 \wedge \neg\varphi_2)$ is equivalent to $\varphi_1 \rightarrow \varphi_2$. Clearly we have $\Omega_{\sigma,\pi}(s) \models \varphi_1$ and $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \rightarrow \varphi_2)$, and hence $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \varphi_2)$. The case of $s \in U_2$ is symmetric. ∎

Example 2 shows that in general we have $U \subsetneq W_{11}$. Given a game graph $G = ((V,E),(V_1,V_2))$ and a subset $V' \subseteq V$ of the states, we write $G \restriction V'$ to denote the subgraph induced by V', that is, $G \restriction V' = ((V', E \cap (V' \times V')), (V_1 \cap V', V_2 \cap V'))$. The following lemma characterizes the set W_{11}.

Lemma 9. $W_{11} = \langle\langle 1,2\rangle\rangle_{G\restriction A}(\varphi_1 \wedge \varphi_2)$.

Proof. Let $s \in \langle\langle 1,2\rangle\rangle_{G\restriction A}(\varphi_1 \wedge \varphi_2)$. The case $s \in U$ is covered by Lemma 8, so let $s \in A \setminus U$. Let (σ,π) be a cooperative strategy profile at the state s, that is, $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \varphi_2)$. Observe that if $t \in A \setminus U$ then $t \notin \langle\langle 1\rangle\rangle_G(\varphi_1)$ and $t \notin \langle\langle 2\rangle\rangle_G(\varphi_2)$. Hence, by the determinacy of the zero-sum games, from every state $t \in A \setminus U$, player 1 (resp. player 2) has a strategy $\overline{\sigma}$ (resp. $\overline{\pi}$) to satisfy the objective $\neg\varphi_2$ (resp. $\neg\varphi_1$) from the state s. We define a pair $(\sigma + \overline{\sigma}, \pi + \overline{\pi})$ of strategies from s as follows. Let $x \in A^*$ be a prefix of a play.

- When the play reaches a state $t \in U$, the players follow their winning retaliating strategies from t. It follows from Lemma 8 that $U \subseteq W_{11}$.
- If $x \in (A \setminus U)^*$, that is, if the play has not yet reached the set U, then player 1 uses the strategy σ and player 2 uses the strategy π. If, however, player 2 deviates from the strategy π, then player 1 switches to the strategy $\overline{\sigma}$ from the first state after the deviation, and symmetrically, if player 1 deviates from σ, then player 2 switches to the $\overline{\pi}$.

It is easy to observe that both strategies $(\sigma + \overline{\sigma})$ and $(\pi + \overline{\pi})$ are retaliating strategies and $\Omega_{\sigma+\overline{\sigma},\pi+\overline{\pi}}(s) \models (\varphi_1 \wedge \varphi_2)$, because $\Omega_{\sigma+\overline{\sigma},\pi+\overline{\pi}}(s) = \Omega_{\sigma,\pi}(s)$. Hence $s \in W_{11}$.

Let $s \notin \langle\langle 1, 2\rangle\rangle_{G\restriction A}(\varphi_1 \wedge \varphi_2)$. Then $s \notin W_{11}$, because for every strategy profile (σ, π) we have either $\Omega_{\sigma,\pi}(s) \models \neg\varphi_1$ or $\Omega_{\sigma,\pi}(s) \models \neg\varphi_2$. ∎

We now define two forms of ω-regular objectives, Rabin and parity objectives. For an infinite path $\Omega = \langle s_0, s_1, s_2, \ldots\rangle$, we define $\mathrm{Inf}(\Omega) = \{ s \in V : s_k = s \text{ for infinitely many } k \geq 0 \}$.

– *Rabin:* We are given a set $\alpha \subseteq 2^V \times 2^V$ of pairs such that $\alpha = \{ (E_1, F_1), (E_2, F_2), \ldots, (E_d, F_d) \}$, where $E_i, F_i \subseteq V$ for all $1 \leq i \leq d$. A *Rabin objective* has the form $\varphi_{Rabin} = \{ \Omega \in V^\omega : \text{there exists } 1 \leq i \leq d \text{ such that } \mathrm{Inf}(\Omega) \cap E_i = \emptyset \text{ and } \mathrm{Inf}(\Omega) \cap F_i \neq \emptyset \}$.
– *Parity:* For $d \in \mathbb{N}$, we write $[d]$ to denote the set $\{0, 1, \ldots, d\}$, and $[d]_+ = \{1, 2, \ldots, d\}$. We are given a function p: $V \to [d]$ that assigns a *priority* $p(s)$ to every state $s \in V$. A *parity (or Rabin chain) objective* has the form $\varphi_P = \{ \Omega \in V^\omega : \min\big(p(\mathrm{Inf}(\Omega))\big) \text{ is even} \}$.

Every ω-regular set can be defined as a parity objective [17]. It follows from Lemma 9 that in order to compute the sets W_{10}, W_{01}, W_{11}, and W_{00}, it suffices to solve two games with conjunctive objectives and a model-checking (1-player) problem for a conjunctive objective. If the objectives φ_1 and φ_2 are ω-regular sets specified as parity objectives, then the conjunctions can be expressed as the complement of a Rabin objective [17]. This gives the following result. (The *size* of a game graph G is $|V| + |E|$).

Theorem 4 (Complexity of computing secure equilibria). *Consider a game graph G of size n, and two Borel objectives φ_1 and φ_2 for the two players.*

– *The four sets W_{10}, W_{01}, W_{11}, and W_{00} can be computed as $W_{10} = \langle\langle 1\rangle\rangle_G(\varphi_1 \wedge \neg\varphi_2)$; $W_{01} = \langle\langle 2\rangle\rangle_G(\varphi_2 \wedge \neg\varphi_1)$; $W_{11} = \langle\langle 1, 2\rangle\rangle_{G\restriction A}(\varphi_1 \wedge \varphi_2)$, where $A = V \setminus (W_{10} \cup W_{01})$; and $W_{00} = V \setminus (W_{10} \cup W_{01} \cup W_{11})$.*
– *If φ_1 and φ_2 are ω-regular objectives specified as LTL formulas, then deciding W_{10}, W_{01}, W_{11}, and W_{00} is 2EXPTIME-complete. The four sets can be computed in time $O\big(n^{2^\ell} \times 2^{2^{\ell \log \ell}}\big)$, where $\ell = |\varphi_1| + |\varphi_2|$ [15].*
– *If φ_1 and φ_2 are parity objectives, then W_{10}, W_{01}, W_{11}, and W_{00} can be decided in co-NP. The four sets can be computed in time $O\big((nd)^{2d}\big)$, where d is the maximal number of priorities in the priority functions for φ_1 and φ_2 [5,4].*

4 $\boldsymbol{\omega}$-Regular Objectives

In this section we consider special cases of graph games, where the two players have reachability, safety, Büchi, co-Büchi, and parity objectives. We fix a game graph G with state space V. Given state sets $R, S, B, C \subseteq V$, these objectives are defined as follows.

1. *Reachability:* $\varphi_R = \{s_0 s_1 \ldots \in V^\omega : \exists k.\ s_k \in R\}$. We refer to R as the *target* set.
2. *Safety:* $\varphi_S = \{s_0 s_1 \ldots \in V^\omega : \forall k.\ s_k \in S\}$. We refer to S as the *safe* set.
3. *Büchi:* $\varphi_B = \{s_0 s_1 \ldots \in V^\omega : \forall k.\ \exists l > k.\ s_l \in B\}$. We refer to B as the *Büchi* set.
4. *co-Büchi:* $\varphi_C = \{s_0 s_1 \ldots \in V^\omega : \exists k.\ \forall l > k.\ s_l \in C\}$. We refer to C as the *co-Büchi* set.

Parity objectives were defined in the previous section. Note that Büchi and co-Büchi objectives are special cases of parity objectives with two priorities: in the Büchi case, take the priority function $p\colon V \to [1]$ such that $p(s) = 0$ if $s \in B$, and $p(s) = 1$ otherwise; in the co-Büchi case, take the priority function $p\colon V \to [2]_+$ such that $p(s) = 2$ if $s \in C$, and $p(s) = 1$ otherwise.

We characterize the memory requirements for strongly winning and retaliating strategies if both players have ω-regular objectives. A retaliation strategy profile (σ, π) is called *winning* at a state $s \in V$ if $\Omega_{\sigma,\pi}(s) \models (\varphi_1 \wedge \varphi_2)$. A strategy σ is a *winning* retaliating strategy for player 1 at state s if there is a strategy π for player 2 such that (σ, π) is a winning retaliation strategy profile at s. Until the end of this section, let φ_R be a reachability objective, φ_S a safety objective, φ_B a Büchi objective, φ_C a co-Büchi objective, and φ_P a parity objective.

Proposition 1 (Conjunctive objectives as parity objectives).

1. $\neg\varphi_R$ is a safety objective and $\neg\varphi_S$ is a reachability objective,
2. $\neg\varphi_C$ is a Büchi objective, and $\neg\varphi_B$ is a co-Büchi objective.
3. $\neg\varphi_P$, $\varphi_S \wedge \varphi_P$, and $\varphi_C \wedge \varphi_P$ are parity objectives.

Proof. A negation of a parity objective with priority function p can be obtained as the parity objective with the priority function $p'(s) = p(s) + 1$. It follows that the negation of a Büchi objective is equivalent to a co-Büchi objective and the negation of a co-Büchi objective is equivalent to a Büchi objective.

If φ_P is a parity objective and φ_D is a safety objective or a co-Büchi objective then the conjunction $\varphi_D \wedge \varphi_P$ is equivalent to a parity objective. For example, the conjunction of a parity objective φ_P and a coBüchi objective φ_D is a parity objective with the following priority function:

$$p'(s) = \begin{cases} 1 & \text{if } s \notin D, \\ p(s) + 2 & \text{if } s \in D. \end{cases}$$

The result for conjunction of parity and safety objective follows from similar construction. ∎

While in zero-sum games played on graphs, memoryless winning strategies exists for all parity objectives [6], this is not the case for non-zero-sum games. The following two theorems give a complete characterization.

Theorem 5. *If player 1 has a strongly winning strategy in a graph game where both players have reachability, safety, Büchi, co-Büchi, or parity objectives φ_1 and φ_2, then player 1 has a memoryless strongly winning strategy if and only if there is a "+" symbol in the corresponding entry of the Table 1.*

Table 1. Strongly winning strategies

		φ_2				
		φ_R	φ_B	φ_C	φ_P	φ_S
φ_1	φ_S	+	+	+	+	+
	φ_C	+	+	+	+	−
	φ_B	+	+	−	−	−
	φ_P	+	+	−	−	−
	φ_R	+	−	−	−	−

Table 2. Winning retaliating strategies

		φ_2				
		φ_R	φ_B	φ_C	φ_P	φ_S
φ_1	φ_S	+	+	+	+	+
	φ_C	+	−	−	−	−
	φ_B	+	−	−	−	−
	φ_P	+	−	−	−	−
	φ_R	+	−	−	−	−

Proof. For player 1, strongly winning a non-zero-sum game with objectives φ_1 and φ_2 is equivalent to winning a zero-sum game with the objective $\varphi_1 \wedge \neg\varphi_2$. Hence by existence of memoryless winning strategies for zero-sum parity games [6] player 1 has memoryless strongly winning strategies if the objective $\varphi_1 \wedge \neg\varphi_2$ is equivalent to a parity objective. Using Proposition 1 it is easy to observe that the objective $\varphi_1 \wedge \neg\varphi_2$ is equivalent to a parity objective for all "+" entries in Table 1, except for safety–reachability, safety–safety, and reachability–reachability games. For these three cases, it is easy to argue that memoryless strongly winning strategies exist. The other "+" entries follow from the existence of memoryless winning strategies for zero-sum parity games [6].

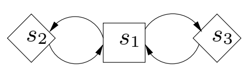

Fig. 3. A counterexample for memoryless strongly winning strategies

We now show that player 1 does not necessarily have a memoryless strongly winning strategy in non-zero-sum games with "−" entries in Table 1. It suffices to give counterexamples for the following four cases: co-Büchi–safety, Büchi–safety, reachability–safety, and Büchi–co-Büchi games. The cases of reachability–Büchi and reachability–co-Büchi games follow from the former two cases, respectively, by symmetry. The cases of Büchi–parity and parity–parity games follow trivially from the Büchi–co-Büchi case, and the case of parity–safety games follows trivially from the Büchi–safety case. The game graph of Fig. 3 serves as a counterexample for all four cases. For all the cases, let $C = S = \{s_1, s_2\}$ and $B = R = \{s_2\}$.

For the co-Büchi–safety case, the player 1 strategy that chooses $s_1 \to s_3$ for the first time and then always chooses $s_1 \to s_2$ is strongly winning at the state s_1, but the two possible memoryless strategies are not strongly winning. For all other cases, the player 1 strategy that alternates between the two moves available at s_1 is strongly winning, but again the two memoryless strategies are not. ∎

Theorem 6. *If player 1 has a winning retaliating strategy in a graph game where both players have reachability, safety, Büchi, co-Büchi, or parity objectives φ_1 and φ_2, then player 1 has a memoryless winning retaliating strategy if and only if there is a "+" symbol in the corresponding entry of the Table 2.*

Proof. First we show that player 1 has memoryless winning retaliating strategies in parity–reachability and safety–parity games. Recall the weakly winning sets $U_1 = W_1 \setminus W_{10}$ and $U_2 = W_2 \setminus W_{01}$, where $W_i = \langle\langle i \rangle\rangle \varphi_i$ for $i \in \{1, 2\}$. In $U_1 \subseteq W_{11}$ player 1 uses her memoryless winning strategy in the zero-sum game with the objective φ_P. In $W_{11} \setminus U_1$ player 1 uses a memoryless strategy that shortens the distance in the game graph to the set U_1. This strategy is a winning retaliating strategy for player 1 in U_1, because it satisfies the objective φ_P. We prove that it is also a winning retaliating strategy for player 1 in $W_{11} \setminus U_1$, that is, satisfaction of the objective φ_R implies satisfaction of the objective φ_P. Observe that $R \cap (W_{11} \setminus U_1) = \emptyset$. Otherwise there would be a state in $W_{11} \setminus U_1$ in which the objective φ_R of player 2 is satisfied and player 2 has a strategy to satisfy $\neg\varphi_P$, and hence the state belongs to W_{01}; this however contradicts $W_{11} \cap W_{01} = \emptyset$. Therefore, as long as a play stays in $W_{11} \setminus U_1$, the objective φ_R cannot be satisfied. On the other hand, if player 2 cooperates with player 1 in reaching U_1, then player 1 plays her memoryless retaliating strategy in U_1. The proof for safety–parity games is similar. There, the key observation is that $W_{11} \setminus U_1 \subseteq S$, where φ_S is the safety objective of player 1.

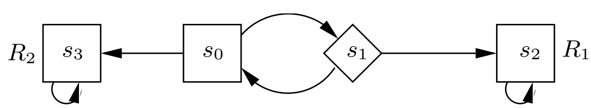

Fig. 4. A counterexample for memoryless winning retaliating strategies

We now argue that player 1 does not have memoryless winning retaliating strategies in games with "−" entries in Table 2. It suffices to give counterexamples for the nine cases that result from co-Büchi, Büchi, or reachability objectives for player 1, and Büchi, co-Büchi, or safety objectives for player 2. The remaining seven cases involving parity objectives follow as corollaries, because Büchi and co-Büchi objectives are special cases of parity objectives. The game graph of Fig. 4 serves as a counterexample for all nine cases: take $C_1 = B_1 = R_1 = \{s_2\}$

and $B_2 = C_2 = S_2 = \{s_0, s_1, s_2\}$, where C_1, B_1, and R_1 are the co-Büchi, Büchi, and reachability objectives of player 1, respectively, and B_2, C_2, and S_2 are the Büchi, co-Büchi, and safety objectives of player 2. It can be verified that in each of the nine games neither of the two memoryless strategies for player 1 is a winning retaliating strategy at the state s_0, but the strategy that first chooses the move $s_0 \to s_1$ and then chooses $s_0 \to s_3$ if player 2 chooses $s_1 \to s_0$, is a winning retaliating strategy for player 1. ∎

Note that if both players have parity objectives, then at all states in W_{00} memoryless retaliation strategy profiles exist. To see this, consider a state $s \in W_{00}$. There are a player 1 strategy $\overline{\sigma}$ and a player 2 strategy $\overline{\pi}$ such that for all strategies σ of player 1 and π of player 2, we have $\Omega_{\sigma,\overline{\pi}}(s) \models \neg\varphi_1$ and $\Omega_{\overline{\sigma},\pi}(s) \models \neg\varphi_2$. The strategy profile $(\overline{\sigma}, \overline{\pi})$ is a retaliation strategy profile. If the objectives φ_1 and φ_2 are both parity objectives, then $\neg\varphi_1$ and $\neg\varphi_2$ are parity objectives as well. Hence there are memoryless strategies $\overline{\sigma}$ and $\overline{\pi}$ that satisfy the above condition.

5 n-Player Games

We generalize the definition of secure equilibria to the case of $n > 2$ players. We show that in n-player games on graphs, in contrast to the 2-player case, there may not be a unique maximal secure equilibrium. The preference ordering $\prec_i$ for player i, where $i \in \{1, \ldots, n\}$, is defined as follows: given two payoff profiles $v = (v_1, \ldots, v_n)$ and $v' = (v'_1, \ldots, v'_n)$, we have $v \prec_i v'$ iff $(v'_i > v_i) \lor (v'_i = v_i \land (\forall j \neq i.\ v'_j \leq v_j) \land (\exists j \neq i.\ v'_j < v_j))$. In other words, player i prefers v' over v iff she gets a greater payoff in v', or (1) she gets equal payoff in v' and v, (2) the payoff of every other player is no more in v' than in v, and (3) there is at least one player who gets a lower payoff in v' than in v. Given a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$, we define the corresponding payoff profile as $v^\sigma = (v_1^\sigma, \ldots, v_n^\sigma)$, where v_i^σ is the payoff for player i when all players choose their strategies from the strategy profile σ. Given a strategy σ'_i for player i, we write (σ_{-i}, σ'_i) for the strategy profile where each player $j \neq i$ plays the strategy σ_j, and player i plays the strategy σ'_i. An n-player strategy profile σ is *Nash equilibrium* if for all players i and all strategies σ'_i of player i, if $\sigma' = (\sigma_{-i}, \sigma'_i)$, then $v_i^{\sigma'} \leq v_i^\sigma$.

Definition 8 (Secure n-player profile). *An n-player strategy profile σ is* secure *if for all players i and $j \neq i$, and for all strategies σ'_j of player j, if $\sigma' = (\sigma_{-j}, \sigma'_j)$, then $(v_j^{\sigma'} \geq v_j^\sigma) \to (v_i^{\sigma'} \geq v_i^\sigma)$.* ∎

Observe that if a secure profile σ is interpreted as a contract between the players, then any unilateral selfish deviation from σ must be cooperative in the following sense: if player j deviates from the contract σ by playing a strategy σ'_j (i.e., the new strategy profile is $\sigma' = (\sigma_{-j}, \sigma'_j)$) which gives her an advantage (i.e., $v_j^{\sigma'} \geq v_j^\sigma$), then every other player $i \neq j$ is not put at a disadvantage if she follows the contract (i.e., $v_i^{\sigma'} \geq v_i^\sigma$). By symmetry, the player j enjoys the same security against unilateral selfish deviations of other players.

Definition 9 (Secure n-player equilibrium). *A n-player strategy profile σ is a* secure equilibrium *if σ is both a Nash equilibrium and secure.* ∎

Similar to Lemma 1 we have the following result.

Lemma 10 (Equivalent characterization). *An n-player strategy profile σ is a* secure equilibrium *iff for all players i, there does not exist a strategy σ_i' of player i such that $\sigma' = (\sigma_{-i}, \sigma_i')$ and $v^{\sigma} \prec_i v^{\sigma'}$.*

We give an example of a 3-player graph game where the maximal secure equilibrium payoff profile is not unique. Recall the game graph from Fig. 3, and consider a 3-player game on this graph where each player has a reachability objective. The target set for player 1 is $\{s_2, s_3\}$; for player 2 it is $\{s_2\}$; and for player 3 it is $\{s_3\}$. In state s_1 player 1 can chose between the two successors s_2 and s_3. If player 1 chooses $s_1 \to s_3$, then the payoff profile is $(1, 0, 1)$, and if player 1 chooses $s_1 \to s_2$, then the payoff profile is $(1, 1, 0)$. Both are secure equilibria and maximal, but incomparable.

6 Conclusion

We considered non-zero-sum graph games with lexicographically ordered objectives for the players in order to capture adversarial external choice, where each player tries to minimize the other player's payoff as long as it does not decrease her own payoff. We showed that these games have a unique maximal equilibrium for all Borel winning conditions. This confirms that secure equilibria provide a good formalization of rational behavior in the context of verifying component-based systems.

Concretely, suppose the two players represent two components of a system with the specifications φ_1 and φ_2, respectively. Classically, component-wise verification would prove that for an initial state s, player 1 can satisfy the objective φ_1 no matter what player 2 does (i.e., $s \in \langle\langle 1 \rangle\rangle \varphi_1$), and player 2 can satisfy the objective φ_2 no matter what player 1 does (i.e., $s \in \langle\langle 2 \rangle\rangle \varphi_2$). Together, these two proof obligations imply that the composite system satisfies both specifications φ_1 and φ_2. The computational gain from this method typically arises from abstracting the opposing player's (i.e., the environment's) moves for each proof obligation. Our framework provides two weaker proof obligations that support the same conclusion. We first show that player 1 can satisfy φ_1 provided that player 2 does not sabotage her ability to satisfy φ_2, that is, we show that $s \in (W_{10} \cup W_{11})$: either player 1 has a strongly winning strategy, or there is a winning pair of retaliation strategies. This condition is strictly weaker than the condition that player 1 has a winning strategy, and therefore it is satisfied by more states. Second, we show the symmetric proof obligation that player 2 can satisfy φ_2 provided that player 1 does not sabotage her ability to satisfy φ_1, that is, $s \in (W_{01} \cup W_{11})$. While they are weaker than their classical counterparts, both new proof obligations together still suffice to establish that $s \in W_{11}$, that is, the composite system satisfies $\varphi_1 \wedge \varphi_2$ assuming that both players behave rationally and follow the winning pair of retaliation strategies.

It should be noted that the other possible lexicographic ordering of objectives captures *cooperative* external choice, where each player tries to *maximize* the other player's payoff as long as it does not decrease her own payoff. However, cooperation does not uniquely determine a preferable behavior: there may be multiple maximal payoff profiles for cooperative external choice, even for reachability objectives. To see this, define $(v_1, v_2) \prec_1^{co} (v_1', v_2')$ iff $(v_1 < v_1') \vee (v_1 = v_1' \wedge v_2 < v_2')$, and $(v_1, v_2) \preceq_1^{co} (v_1, , v_2')$ iff $(v_1, v_2) \prec_1^{co} (v_1', v_2')$ or $(v_1, v_2) = (v_1', v_2')$. A symmetric definition yields $\preceq_2^{co}$. A *cooperative equilibrium* is a Nash equilibrium with respect to the precedence orderings $\preceq_1^{co}$ and $\preceq_2^{co}$ on payoff profiles. Consider the game shown in Fig. 4, where each player has a reachability objective. The target for player 1 is s_2, and the target for player 2 is s_3. The possible cooperative equilibria at state s_0 are as follows: player 1 chooses $s_0 \to s_1$ and player 2 chooses $s_1 \to s_2$, or player 1 chooses $s_0 \to s_3$ and player 2 chooses $s_1 \to s_0$. The former equilibrium has the payoff profile $(1, 0)$, and the latter has the payoff profile $(0, 1)$. These are the only cooperative equilibria and, therefore, the maximal payoff profile for cooperative equilibria is not unique.

Acknowledgment. We thank Christos Papadimitriou for helpful discussions regarding the formalization of rational behavior in game theory.

References

1. M. Abadi and L. Lamport. Conjoining specifications. *ACM Transactions on Programming Languages and Systems*, 17:507–534, 1995.
2. R. Alur and T.A. Henzinger. Reactive modules. In *Formal Methods in System Design*, 15:7–48, 1999.
3. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
4. S. Dziembowski, M. Jurdziński, and I. Walukiewicz. How much memory is needed to win infinite games? In *Logic in Computer Science* (LICS), pages 99–110. IEEE Computer Society Press, 1997.
5. E.A. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *Foundations of Computer Science* (FOCS), pages 328–337. IEEE Computer Society Press, 1988.
6. E.A. Emerson and C. Jutla. Tree automata, μ-calculus, and determinacy. In *Foundations of Computer Science* (FOCS), pages 368–377. IEEE Computer Society Press, 1991.
7. J.F. Nash Jr. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.
8. A. Kechris. *Classical Descriptive Set Theory.* Springer-Verlag, 1995.
9. D.M. Kreps. *A Course in Microeconomic Theory.* Princeton University Press, 1990.
10. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification.* Springer-Verlag, 1992.
11. D.A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.
12. D.A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63:1565–1581, 1998.

13. K. Namjoshi N. Amla, E.A. Emerson and R. Trefler. Abstract patterns for compositional reasoning. In *Concurrency Theory* (CONCUR), LNCS 2761, pages 423–448. Springer-Verlag, 2003.
14. G. Owen. *Game Theory.* Academic Press, 1995.
15. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Principles of Programming Languages* (POPL), pages 179–190. ACM Press, 1989.
16. W. Thomas. On the synthesis of strategies in infinite games. In *Symposium on Theoretical Aspects of Computer Science* (STACS), LNCS 900, pages 1–13. Springer-Verlag, 1995.
17. W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, eds., *Handbook of Formal Languages*, volume 3, pages 389–455. Springer-Verlag, 1997.