

Equilibria of Concurrent Games on Event Structures

Julian Gutierrez Michael Wooldridge

Department of Computer Science
University of Oxford

Abstract

Event structures form a canonical model of concurrent behaviour which has a natural game-theoretic interpretation. This game-based interpretation was initially given for zero-sum concurrent games. This paper studies an extension of such games on event structures to include a much wider class of game types and solution concepts. The extension permits modelling scenarios where, for instance, co-operation or independent goal-driven behaviour of computer agents is desired. Specifically, we will define non-zero-sum games on event structures, and give full characterisations—existence and completeness results—of the kinds of games, payoff sets, and strategies for which Nash equilibria and subgame perfect Nash equilibria always exist. The game semantics of various logics and systems are outlined to illustrate the power of this framework.

Categories and Subject Descriptors Theory of Computation [Logics and Meanings of Programs]: Semantics.

General Terms Theory.

Keywords Event structures, Concurrent games, Nash equilibria.

1. Introduction

Game theoretic concepts have a venerable history in logic and computer science. For example, game theoretic concepts have been applied to great effect in areas such as satisfiability, model checking, equivalence checking, synthesis and realizability, and control, amongst others (see, *e.g.*, [5, 8, 16, 18, 39] for surveys). A common technique is to formulate a decision problem on a logical or mathematical structure S as a two-person zero-sum game, so that a particular player has a winning strategy in the game iff the structure S is a positive instance of the decision problem. In our work, we are also interested in using game theoretic techniques for the analysis of computational systems, but our motivation is rather different, and as a consequence, the game theoretic concepts of relevance are also different. Our aim is to develop the theory required to understand the behaviour of concurrent systems, under the assumption that the computational components of these systems have their own private goals, and are assumed to act *strategically* in the pursuit of these goals. We do *not* assume that the goals of different system components are in conflict: while this may be the case, it may also

be the case that goals are mutually consistent. Such systems cannot be modelled as zero-sum win/lose games, of the type used to such great effect for other purposes in logic and computer science. Instead, we use the more widely applicable model of *non-zero-sum games*, and the associated analytical concepts of *Nash equilibrium* and its many variants. Thus, we consider concurrent systems in which the components of that system have individual goals, and are assumed to act strategically in pursuit of these goals; and the basic question we ask of such a computational system will be:

What are, if any, the Nash equilibria of the system?

i.e., what computations will occur if the players of that concurrent system (are designed to) act rationally and strategically?

Similar ideas have been studied in several areas of computer science: for example, equilibrium analysis has been a major theme in the algorithmic game theory community for over a decade [31]. In logic and verification these ideas have attracted somewhat less attention [8, 10, 19]. However, these latter works have made use of games on graphs, implicitly corresponding either to sequential computations or to interleaved semantics for concurrency.

The present paper is concerned with equilibrium analysis in the context of concurrent systems with partial order behaviour, *i.e.*, with causal or *noninterleaving* semantics [29]. In such a case the problem at hand may be rather different, both technically and conceptually. Indeed, very little is known about non-zero-sum games for concurrent systems with causal or noninterleaving semantics. This represents a logical and computational problem when analysing concurrent systems at a *fundamental semantic level*, since they then have to be interpreted under their semantically less informative interleaving interpretations—a situation that is problematic even in the zero-sum setting [2]. We want to avoid a semantically undesirable translation from noninterleaving to interleaving models when analysing concurrent systems modelled as games, and be able to work, use, and ultimately fully understand non-zero-sum concurrent games played *directly* on the original noninterleaving models where they were defined. This is, in essence, the problem we want to study and main motivation.

We study non-zero-sum games in the general context of event structures [29], a computationally powerful, yet mathematically simple model of concurrency. Event structures are the concurrency analogue of trees. While event structures have natural game-based interpretations [13, 21, 33], the games studied in [13, 21, 33] are zero-sum, and hence inappropriate for our purposes, as discussed above. In this paper we broaden the class of game-theoretic concepts that can be applied to the analysis of logics and systems with event structure semantics. Specifically, we extend [13, 21] to allow Borel *non-zero-sum* payoff sets and define *equilibrium computations* in this more general setting. In particular, we formulate and investigate the fundamental notions of *Nash equilibrium* as well as *subgame perfect Nash equilibrium* for event structures—and hence for concurrent systems with explicit partial order behaviour.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.
<http://dx.doi.org/10.1145/2603088.2603145>

There are very good reasons to study the equilibria of concurrent systems with causal semantics—*i.e.*, with partial order behaviour. The equilibria of a multi-agent system characterise those situations (computations) where no agent (computer program) can unilaterally improve its behaviour [34]. This feature can be used as a formal quality measure in the design of models and verification of computer systems. This is particularly relevant in the context of concurrent systems with causal models since they underpin powerful approaches to formal verification (*e.g.*, partial-order reduction methods [17] and unfolding techniques [15]) as well as semantically finer representations/models of complex computer interactions. It is thus computationally, logically, and semantically convenient to be able to study the equilibria of concurrent systems directly over their potentially more compact, yet semantically more informative, noninterleaving representations. Equilibria as solution concepts for non-zero-sum games are therefore *the* guideline for the better construction of game models of systems [34]. In this paper we study this problem and give a number of positive results for a very general class of event structures—and through them for computer systems and logics with an event structure semantics.

The rest of the paper is organised as follows. In Section 3 we introduce *non-zero-sum Borel* payoff sets for concurrent games on event structures and their associated equilibrium solution concepts. In Section 4 it is shown that two structural conditions on games, called race-freedom and bounded-concurrency, are *necessary* for the existence of equilibria, and in Section 5 we show that such conditions are also *sufficient*—thus implying that no more general computational framework can be established within the model of event structures. Then, in Section 6, the main result of Section 5 is extended to subgame perfect Nash equilibrium computations—a refinement of Nash equilibria which takes into account the dynamic nature of systems where players may *repeatedly* interact with each other. This paper constitutes the first contribution of its kind towards a *theory of equilibria for noninterleaving concurrency*.

2. Preliminaries

Non-zero-sum games and rational behaviour The behaviour of a program P can be modelled as a two-player game between P and an arbitrary environment O (*cf.* [1, 24]). Since O can be any environment, one usually assumes that it can behave in any possible way, namely by providing P with any allowed input in the game. This may be a too specific assumption since one then also has to assume that if P is a process whose task or goal is γ then the goal of O may be to prevent P from achieving γ , that is, to behave in such a way that the goal $\neg\gamma$ is achieved instead. Modelled as games, this situation leads to the construction of a *zero-sum* game, that is, a (two-player, perfect-information) game where the goals of each player are in conflict—the system O is assumed to be a hostile environment. The main decision problem at hand, in game-theoretic terms, then becomes the computation of a winning strategy for one of the players. A winning strategy for P delivers a computation plan to ensure that γ is achieved, whereas one for O provides a way in which O can prevent P from achieving its goal, that is, a way to ensure that the goal $\neg\gamma$ is achieved instead.

Frequently, however, this is not a realistic computational scenario (*e.g.* see [10, 11, 19]) since O itself may be another program that is meant to carry out a particular task, which is not necessarily to prevent P from doing its own. In this more general setting, the interaction between P and O is better represented by a *non-zero-sum* game where the goals γ_P of P and γ_O of O are not necessarily in conflict, that is, it is not necessarily the case that $\gamma_O = \neg\gamma_P$. In this case, the main problem when modelling the system as a two-player game is no longer finding a winning strategy as, in general, it may not exist. Instead, what one wants is to find a pair of strategies, one

for each player, such that they form an equilibrium: a computation plan from which neither P nor O wants to unilaterally deviate.

To give an example, take the concurrent system depicted in Figure 1, where process P is a computer program whose goal γ_P is to be able to execute the event B infinitely often. Using LTL as a specification language, we can say that the goal of P is $\gamma_P = \square\Diamond B$. Such a goal cannot be ensured in the composed system $P \parallel O$ since O controls A and, hence, could prevent P from achieving γ_P . However, even though we may not know the implementation of O , we may know that it is, for instance, a simple scheduler or service provider whose goal is to grant access A whenever, and only if, a request R is made. Such a situation can be modelled by saying that O 's goal is $\gamma_O = \square(R \Leftrightarrow \Diamond A)$. Clearly, $\gamma_O \neq \neg\gamma_P$.




Figure 1: An event (square) is enabled and can happen whenever there is a token (black dot) in each place (circle) it depends on. For instance, event R is enabled and when it happens all tokens are consumed by event R and new tokens are produced as shown above—illustrated with \Downarrow_R : one token for event A and one token for the unknown part of O . Program/Process P has control over events B and R ; environment O controls event A .

If we assume that O behaves “rationally,” meaning that it is a computer system which intends to satisfy its temporal logic specification, rather than assuming that it can have any behaviour, then we know that P can achieve its goal γ_P . In other words, instead of being interested in ensuring that the set of *all* computations of the composed system $P \parallel O$ satisfies the specification γ_P , one is *only* interested in those “equilibrium computations” (defined later) where P and O have no desire to unilaterally deviate. Indeed, in the example above, even though process P cannot ensure on its own that its goal γ_P can always be achieved, it can actually do so when interacting with any environment O whose implementation satisfies the specification γ_O . Figure 1 shows a situation where *non-zero-sum games should be used rather than zero-sum games*. In fact, with respect to our definitions of equilibrium (in Section 3) every rational implementation of γ_O must contain an equilibrium computation that satisfies both γ_P and γ_O .

In the remainder of this section we will formally introduce the model of event structures as well as concurrent games played on them. In particular, the fact that event structures are the concurrency analogue of trees, should help understand that just as *sequential* games can be represented by *trees*, so *concurrent* games can be represented by *event structures*, where plays in this much more general setting determine *partial* instead of *total* orders of moves.

2.1 Event structures

An *event structure* [29] is a triple (E, \leq, Con) , consisting of a set E , of *events* which are partially ordered by \leq , the *causal dependency relation*, and a nonempty *consistency relation* Con consisting of finite subsets of E , which satisfy four axioms:

- (i) $\{e' \mid e' \leq e\}$ is finite for all $e \in E$,
- (ii) $\{e\} \in \text{Con}$ for all $e \in E$,
- (iii) $Y \subseteq X \in \text{Con} \implies Y \in \text{Con}$, and
- (iv) $X \in \text{Con} \ \& \ e \leq e' \in X \implies X \cup \{e\} \in \text{Con}$.

The *configurations*, denoted by $\mathcal{C}(E)$, of an event structure E consist of those subsets $x \subseteq E$ which are both

- *consistent*: $\forall X \subseteq x. X \text{ is finite} \implies X \in \text{Con}$; and
- *down-closed*: $\forall e, e'. e' \leq e \in x \implies e' \in x$.

Configurations can be finite or infinite. We write $\mathcal{C}^\infty(E)$ for the *infinite* configurations of E and $\mathcal{C}^\omega(E)$ for the *finite* ones. Two events, say e and e' , which are both consistent and incomparable with respect to causal dependency are seen as *concurrent*, and denoted by *co*, *i.e.*, $e \text{ co } e'$. In games the relation of *immediate* dependency $e \rightarrow e'$, meaning e and e' are distinct with $e \leq e'$ and no event in between, is important. For $X \subseteq E$ we write $[X]$ for $\{e \in E \mid \exists e' \in X. e \leq e'\}$, the down-closure of X ; note if $X \in \text{Con}$ then $[X] \in \text{Con}$. We use $x \text{--}c y$ to mean y covers x in $\mathcal{C}(E)$, *i.e.* $x \subset y$ with nothing in between, and $x \text{--}^e c y$ to mean $x \cup \{e\} = y$ for $x, y \in \mathcal{C}(E)$ and event $e \notin x$. We write $x \text{--}^e c$ to express that event e is enabled at configuration x , when $x \text{--}^e c y$ for some y . Given an event structure E and a subset $R \subseteq E$ of its events, the *restriction* $E \upharpoonright R$ is the event structure comprising events $\{e \in E \mid [e] \subseteq R\}$ with causal dependency and consistency inherited from E .

A (*partial*) *map* of event structures is a partial function on events $f : E \rightarrow E'$ such that for all $x \in \mathcal{C}(E)$ its direct image $f x \in \mathcal{C}(E')$ and if $e_1, e_2 \in x$ and $f(e_1) = f(e_2)$ (with both defined) then $e_1 = e_2$ —*i.e.*, injective on configurations. The map expresses how the occurrence of an event e in E induces the coincident occurrence of the event $f(e)$ in E' whenever it is defined. Partial maps of event structures compose as partial functions, with identity maps given by identity functions; the map is *total* if f is total. A total map is *strategic* when it respects causal dependency within configurations x , *i.e.*, for all $e, e' \in x$ if $f(e) \rightarrow f(e')$ then $e \rightarrow e'$.

The category of event structures is rich in useful constructions on concurrent processes. For instance, it has products and pullbacks (both forms of synchronised composition) as well as coproducts (nondeterministic sums). In particular, pullbacks will be used to define the composition of *strategies* on event structures. Most formal definitions and results—*e.g.*, existence of the constructions above mentioned—can be found in [29, 40], and references therein.

2.2 Concurrent games on event structures

As traditionally defined, a concurrent game on an event structure is played by two players, ‘Player’ and ‘Opponent’, who interact with each other by choosing, independently and asynchronously, events they control. Positive (+) events are associated with Player and negative (−) events are associated with Opponent. Games on event structures are asynchronous, nondeterministic, and round-free.

We shall informally describe the three above features in more detail. *Asynchrony* means that both players can make multiple and independent choices of events without necessarily waiting for the other to play. *Nondeterminism* is explicitly given by the event structures model: in these games players’ strategies are modelled by (maps of) event structures, which will allow for nondeterministic strategic behaviour. *Round-freeness* is a consequence of the poten-

tial unbounded asynchronous behaviour that event structures allow; it is not known, at any point in time while playing the game, when a player will make a move, in case it does so at all—in fact, no player is forced to make a move in these games.

Moreover, in a game on an event structure each player has a goal, a set of winning plays they want to achieve—*i.e.* their goal is to enforce a run (of the event structure) that is in their winning set. In this paper, we consider *Borel* sets, that is, sets constructed from the open sets in the Scott topology [30] associated with event structures and only using the operations of countable union, countable intersection, and relative complement.

Event structures with polarity Both a concurrent game and a strategy in a game are represented as event structures with polarity, comprising an event structure E together with a polarity function $\text{pol} : E \rightarrow \{+, -\}$ ascribing a polarity + (Player) or − (Opponent) to its events; the events correspond to moves. Maps of event structures with polarity, are maps of event structures which preserve polarities. An event structure with polarity E is *deterministic* iff

$$\forall X \subseteq_{\text{fin}} E. \text{Neg}[X] \in \text{Con}_E \implies X \in \text{Con}_E,$$

where $\text{Neg}[X] =_{\text{def}} \{e' \in X \mid \text{pol}(e') = - \ \& \ \exists e \in E. e' \leq e\}$. We write $\text{Pos}[X]$ when $\text{pol}(e') = +$. We write x^+ for the set of positive (+ve) events in x and write x^- for the set of negative (−ve) ones. And given two sets of events x and y , we write $x \text{ } ^+ \text{ } y$ to express that $x \subset y$ and $\text{pol}(y \setminus x) = \{+\}$; similarly, we write $x \text{ } ^- \text{ } y$ iff $x \subset y$ and $\text{pol}(y \setminus x) = \{-\}$.

Event structures with polarity support two key operations: The *dual*, E^\perp , of an event structure with polarity E comprises the same underlying event structure E but with a reversal of polarities. The *simple parallel composition* $E \parallel E'$ forms the disjoint juxtaposition of E, E' , two event structures with polarity; a finite subset of events is consistent if its intersection with each component is consistent.

Concurrent games and strategies Let A be an event structure with polarity, thought of as a game; its events stand for the possible moves of Player and Opponent and its causal dependency and consistency relations the constraints imposed by the game. A strategy represents a nondeterministic play of the game—all its moves are moves allowed by the game and obey the constraints of the game.

A *strategy* in A is a strategic map $\sigma : S \rightarrow A$ from an event structure with polarity S , which is both *receptive* and *innocent*. Receptivity ensures an openness to all possible moves of Opponent. Innocence, on the other hand, restricts the behaviour of Player; Player may only introduce new relations of immediate causality of the form $\ominus \rightarrow \oplus$ beyond those imposed by the game. Formally:

Receptivity: A map σ is *receptive* iff

$$\sigma x \text{--}^a c \ \& \ \text{pol}(a) = - \implies \exists! s \in S. x \text{--}^s c \ \& \ \sigma(s) = a.$$


Innocence: A map σ is *innocent* iff

$$s \rightarrow s' \ \& \ (\text{pol}(s) = + \text{ or } \text{pol}(s') = -) \text{ then } \sigma(s) \rightarrow \sigma(s').$$

Informally, a strategy $\sigma : S \rightarrow A$ in an event structure A tells Player how to play the game, *i.e.*, how to play in A . The event structure S indicates (i) which events in A are played, (ii) which ones are not played, and (iii) whether an event in A is played provided that some others have been already played—via causal dependencies. We say that a strategy $\sigma : S \rightarrow A$ is deterministic if S is deterministic.


Because we will later on consider winning and losing plays for each player, we can define winning strategies and determined games in the usual way. Informally, a winning strategy is a strategy that ensures a winning outcome for the player that uses it, and a determined game is a game where winning strategies always exist. A determinacy result for certain concurrent games on event structures with Borel winning sets of configurations is known [21].

EXAMPLE 1. Let $\sigma_i : S_i \rightarrow A$ be a strategy in $A = \oplus \text{co} \ominus$



The event structures at the top define the behaviour of Player; at the bottom, the game A is depicted. Dotted lines are the maps. W.r.t. Player's behaviour, σ_0 does nothing, σ_1 plays \oplus regardless of the behaviour of Opponent, and σ_2 awaits and plays \oplus only if Opponent plays \ominus (by adding the causal relation $\ominus \rightarrow \oplus$ in S_2). □

Composing strategies The composition of concurrent strategies can be defined via pullbacks (see, e.g., [29, 40] for operations in the category of event structures). Suppose $\sigma : S \rightarrow A$ is a strategy in a concurrent game A . A counter-strategy is a strategy of Opponent, so a strategy $\tau : T \rightarrow A^\perp$ in the dual game. Ignoring polarities, we have total maps of event structures $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A$. Form their pullback,



to obtain the event structure P resulting from the interaction $\tau \circ \sigma$ of σ and τ . Because σ or τ may be nondeterministic there can be more than one maximal configuration z in $\mathcal{C}(P)$. A maximal z maps to a configuration $\sigma \Pi_1 z = \tau \Pi_2 z$ in $\mathcal{C}(A)$. Define the *results* of playing σ against τ to be

$$\langle \sigma, \tau \rangle =_{\text{def}} \{ \sigma \Pi_1 z \mid z \text{ is maximal in } \mathcal{C}(P) \}.$$

EXAMPLE 2 (Results and composition—from [21]). Consider the game in Example 1 again. As there are three analogous counter-strategies $\tau_j : T_j \rightarrow A^\perp$, $j = 0, 1, 2$, for Opponent, the results in $\langle \sigma_i, \tau_j \rangle$ of playing each σ_i against each τ_j are as follows: $\{\emptyset\}$ if $i \in \{0, 2\}$ & $j \in \{0, 2\}$, $\{\{\oplus\}\}$ if $i = 1$ & $j = 0$, $\{\{\ominus\}\}$ if $i = 0$ & $j = 1$, and $\{\{\oplus, \ominus\}\}$ if $i = 1$ & $j = 1$. □

3. Non-zero-sum concurrent games

A concurrent game with non-zero-sum winning conditions (a non-zero-sum game hereafter) is a triple $G = (A, W_+, W_-)$ where A is an event structure with polarity, $W_+ \subseteq \mathcal{C}(A)$ and $W_- \subseteq \mathcal{C}(A)$ consists of the *winning configurations* for Player and Opponent, respectively. Define the *losing conditions* to be $L_+ = \mathcal{C}(A) \setminus W_+$ and $L_- = \mathcal{C}(A) \setminus W_-$. The dual G^\perp of a non-zero-sum game $G = (A, W_+, W_-)$ is defined to be $G^\perp = (A^\perp, W_-, W_+)$, a game where the roles of Player and Opponent are reversed. The *inverse game* G^{-1} of $G = (A, W_+, W_-)$ is $G^{-1} = (A, L_-, L_+)$, a game where the *goal* (given by the winning configurations) of each player is to prevent the other to win w.r.t. the original game G .

A strategy in G is a strategy in A . A strategy in G is regarded as *winning (for Player)* if it always prescribes Player moves to end up in a winning configuration, i.e. in W_+ , no matter what Opponent does. Formally, a strategy $\sigma : S \rightarrow A$ in G is *winning (for Player)* if $\sigma x \in W_+$ for all \oplus -maximal configurations $x \in \mathcal{C}(S)$ —a configuration x is \oplus -maximal if whenever $x \xrightarrow{s} c$ then the event s has $-ve$ (negative) polarity. We write $\text{max}_+(x)$ if x is \oplus -maximal and $\overline{\text{max}}_+(x)$ if is not; similarly we write $\text{max}_-(x)$ and $\overline{\text{max}}_-(x)$ for \ominus -maximal configurations. Equivalently, a strategy σ for Player is winning if when played against any counter-strategy τ of Opponent, the final result is a win for Player. It is known that a strategy σ is winning for Player iff all the results of the interaction $\langle \sigma, \tau \rangle$ lie within W_+ , for any strategy τ of Opponent; see [13]. The analogous definitions and results apply to Opponent.

Note that we are considering two-player, non-deterministic, pure strategy (i.e., not randomised) games. Further extensions (e.g., to n -player games, mixed strategies, etc.) are possible if the categorical definition of *composition via pullbacks* is in place.

Nash Equilibrium Player and Opponent are associated with strict preference relations, \succ_+ and \succ_- respectively, over the results of a game. Each preference relation is dichotomous in the following sense: players strictly prefer results that are winning over losing ones, and are indifferent otherwise, i.e. for all results $y, y' \in \mathcal{C}(A)$,

$$y \succ_+ y' \iff (y \in W_+ \ \& \ y' \in L_+).$$

Moreover, we write $y \succeq_+ y'$ to mean that it is not the case that $y' \succ_+ y$. Analogous definitions and notations are made and used for Opponent (simply by replacing $+$ with $-$).

A *strategy profile* (σ, τ) is defined to be a (*pure strategy*) *Nash equilibrium* iff for all results $y \in \langle \sigma, \tau \rangle$, strategies σ' for Player, strategies τ' for Opponent, results $y' \in \langle \sigma', \tau \rangle$, and results $y'' \in \langle \sigma, \tau' \rangle$ we have both $y \succeq_+ y'$ and $y \succeq_- y''$.

Let $NE(G)$ be the set of Nash equilibria of G . We say a class of games \mathcal{G} is *complete* iff $NE(G) \neq \emptyset$ for all G in \mathcal{G} . Note the definition of (pure strategy) Nash equilibrium in classical game theory is with respect to strategy profiles that induce a unique outcome, rather than multiple ones as in our more general nondeterministic setting; the two definitions coincide for deterministic games.

It is important to note that the standard definition of Nash equilibrium is either for deterministic or probabilistic strategy profiles. *Nondeterministic* notions have been studied only recently (in computer science) and have no standard counterpart in the classical game theory literature.


The binary character of preferences over outcomes with respect to \succeq_+ and \succeq_- implies the following key fact.

PROPOSITION 3. The strategy profile (σ, τ) is a Nash equilibrium strategy profile iff one of the following holds:

1. $\langle \sigma, \tau \rangle \subseteq W_+ \cap W_-$, or
2. $\langle \sigma, \tau \rangle \subseteq W_+$ and $\langle \sigma, \tau' \rangle \subseteq L_-$ for all τ' , or
3. $\langle \sigma, \tau \rangle \subseteq W_-$ and $\langle \sigma', \tau \rangle \subseteq L_+$ for all σ' , or
4. $\langle \sigma, \tau \rangle \subseteq L_+ \cap L_-$ and $\langle \sigma', \tau \rangle \subseteq L_+$ and $\langle \sigma, \tau' \rangle \subseteq L_-$, for all σ' and τ' .

The proposition above is proved by analysing the outcomes of each strategy profile. Case 1 is a Nash equilibrium because Player and Opponent have their goals achieved; cases 2 and 3 imply the existence of a winning strategy—and any strategy profile containing a winning strategy is a Nash equilibrium; case 4 imply the existence of winning strategies in the inverse game, which induce a Nash equilibrium in the original game.

EXAMPLE 4 (Nash equilibria). Let G be the game (A, W_+, W_-) , with A as depicted below and $W_+ = W_- = \{\{\oplus_R, \ominus_R\}\}$. There are two Nash equilibria where players make a move, namely (σ_L, τ_L) and (σ_R, τ_R) where, informally, σ_L plays \oplus_L and τ_L plays \ominus_L and where σ_R plays \oplus_R and τ_R plays \ominus_R . Wiggly lines mean conflict between events, i.e. inconsistency.



In the non-zero-sum game above no player can unilaterally ensure to end in the desired winning configuration $\{\oplus_R, \ominus_R\}$. However, if they were to act rationally, and in the furtherance of their goals, cooperation can be rationally sustained—by choosing the appropriate Nash equilibrium strategy profile—so that both of their goals are achieved. But not all games in event structures have Nash equilibria. In the following section we will characterise those games where cooperation is irremediably impossible; even more, those scenarios where both players are powerful enough to ensure that the other cannot achieve its goal in some cases.

4. Games without equilibria

Let us start with an example. The next example shows a game on an event structure where no Nash equilibrium exists.

EXAMPLE 5. Let the game G be (A, W_+, W_-) , with $A = \oplus \sim \ominus$, $W_+ = \{\{\oplus\}\}$, and $W_- = \{\{\ominus\}\}$. There are two strategies for Player and two for Opponent, namely the strategy σ_1 (resp. τ_1) that plays \oplus (resp. \ominus) and the strategy σ_2 (resp. τ_2) that does not play such an event. The results of playing σ_i against τ_j , with $i, j \in \{1, 2\}$, are: $\{\emptyset\}$ if $i = j = 2$, $\{\{\oplus\}\}$ if $i = 1$ & $j = 2$, $\{\{\ominus\}\}$ if $i = 2$ & $j = 1$, and $\{\{\oplus\}, \{\ominus\}\}$ if $i = 1$ & $j = 1$. Then, the possibility of switching to play either σ_1 or τ_1 makes every (σ_i, τ_j) not a Nash equilibrium. Observe that in contrast to the games studied in Economics, allowing mixed strategies would not help here: in this game there is no Nash equilibrium in mixed strategies either—a good reason to only consider pure strategies. \square

Once the concept of Nash equilibrium has been introduced, the most natural and important question to answer is:¹

Given: Game G .

NONEMPTINESS: Is it the case that $NE(G) \neq \emptyset$?

Our first incompleteness result states that whenever the next structural property, called *race-freedom*, is not satisfied by a game A , then there are winning conditions W_+ and W_- such that the set $NE(A, W_+, W_-)$ is empty, i.e. such that the class of games containing (A, W_+, W_-) is incomplete. We say that a game A is race-free if and only if for all configurations $x \in \mathcal{C}(A)$ and events $a, a' \in A$ we have that

$$x \xrightarrow{a} \text{c} \ \& \ x \xrightarrow{a'} \text{c} \ \& \ \text{pol}(a) \neq \text{pol}(a') \implies x \cup \{a, a'\} \in \mathcal{C}(A). \quad (\mathbf{Race - free})$$

Race-freedom ensures a natural property of concurrent games, namely that each player makes its choices independently of the other—i.e. in a concurrent way.

THEOREM 6. Let A be a game on an event structure that is not race-free. Then there are W_+ and W_- such that the set $NE(G) = \emptyset$, with $G = (A, W_+, W_-)$.

To prove Theorem 6, note that since A is not race-free we can construct a zero-sum game G that is not determined [21]. We can, thus, suppose that (σ, τ) is a Nash equilibrium strategy profile in order to get a contradiction. By Proposition 3 we know that one of cases 1–4 should hold. However, cases 1 and 4 are impossible because G is zero-sum, and cases 2 and 3 are impossible because G is not determined—which leads to the desired contradiction.

Using virtually the same proof technique as in Theorem 6 we can also show that if a concurrent game A does not satisfy a structural property called *bounded-concurrency* then, again, there will be winning conditions W_+ and W_- such that the set of Nash Equilibria $NE(A, W_+, W_-)$ is empty. We say that a game A is bounded-concurrent if and only if for all configurations $x \in \mathcal{C}(A)$ and events $a \in x$ we have that

$$\{a' \in x \mid a \text{ co } a' \ \& \ \text{pol}(a) \neq \text{pol}(a')\} \text{ is finite.} \quad (\mathbf{Bounded - concurrent})$$

The second incompleteness result then states the following.

THEOREM 7. Let A be a game on an event structure that is not bounded-concurrent. Then there are W_+ and W_- such that the set $NE(G) = \emptyset$, with $G = (A, W_+, W_-)$.

¹In 1951 John Nash answered this question positively for games with a finite number of players where each player can probabilistically choose from a finite set of pure strategies—provided the games are restricted to those with a space of strategies that is non-empty, convex, and compact [28].

The interpretation of Theorems 6 and 7 is simple: these two theorems simply imply that race-freedom and bounded-concurrency are *necessary* conditions for the existence of Nash equilibria. To be more precise, note that given a game (A, W_+, W_-) , necessity of these two structural conditions means that if either of them does not hold in A , then we can find some W'_+ and W'_- (which may well be *different* from W_+ and W_-) such that the game (A, W'_+, W'_-) does not have a Nash equilibrium. In the next section we will show that both structural properties are, moreover, *sufficient* with respect to Borel sets of winning configurations/plays, that is, they can be used to ensure the existence of equilibria provided that (i) A is race-free and bounded-concurrent and (ii) both W_+ and W_- are Borel sets.

Remark Both race-freedom and bounded-concurrency are satisfied by the best-known concurrent games in the computer science literature, e.g., by games on concurrent game structures [3] or by games on graphs and trees [20], even in multi-player concurrent settings. Thus, there are no semantic restrictions with respect to asynchronous games played on such kinds of structures. In Section 7, encodings of some of such games are described in order to illustrate how both race-freedom and bounded-concurrency arise in those cases. Informally, race-freedom is satisfied because in the games mentioned above there are rounds, wherein the moves available to each player cannot be disabled in the same round by other players—implying race-freedom. On the other hand, bounded-concurrency is satisfied because at each round, say round i , every player has only finitely many available moves and since the set of players is finite, bounded-concurrency within rounds immediately follows. Moreover, no event in round $i + 1$ can be played (i.e., is enabled) before round i finishes—thus, implying bounded-concurrency in our true-concurrency semantic framework.

Note that the two structural properties studied in this section are interesting only if the model or structure where the game is played—the arena game—has a *noninterleaving* (sometimes called ‘truly concurrent’) semantics [29]; for instance, in models of concurrency such as Petri nets, asynchronous transition systems, Mazurkiewicz traces, or, of course, event structures.

5. Existence of equilibria

In this section we present two results that give a full semantic characterisation of the equilibrium computations of concurrent games on event structures. The first one states, as noted before, that if a concurrent game is race-free and bounded-concurrent, then such a game always has a Nash equilibrium provided its sets of winning configurations are Borel. The second one states, moreover, that for each strategy profile that forms a Nash equilibrium, there is an equivalent strategy profile within a class of nondeterministic concurrent strategies which we call defensive [22]. We now define the classes of games and strategies considered in this section.

Games

A concurrent game with non-zero-sum Borel winning conditions (a Borel non-zero-sum game) is a non-zero-sum game $G = (A, W_+, W_-)$ where W_+ and W_- are Borel subsets of the set of configurations $\mathcal{C}(A)$. Non-zero-sum concurrent games on event structures that are race-free and bounded-concurrent enjoy several useful properties, most notably the one to be presented next, whose proof critically depends on the games being both race-free and bounded-concurrent. Before stating the result, let us first introduce some notations. Let $L_+^\infty = L_+ \cap \mathcal{C}^\infty(A)$ and $W_+^\infty = W_+ \cap \mathcal{C}^\infty(A)$, and recall that c^+ denotes strict inclusion, provided that if $y \subset^+ y'$ then all events in $y' \setminus y$ have positive polarity.

With the definitions and notations given above, we can now present the following simple technical result:

LEMMA 8. Let (A, W_+, W_-) be a non-zero-sum concurrent game that is race-free and bounded-concurrent. Then, for all infinite configurations $y, y' \in \mathcal{C}^\omega(A)$, if $y \in L_+^\infty$ and $y' \in W_+^\infty$ and $y \subset^+ y'$, then there is a necessarily finite and \ominus -maximal configuration $y'' = [\text{Neg}[y]] \in \mathcal{C}^\omega(A)$ such that both $y'' \subset^+ y$ and $y'' \subset^+ y'$. Similarly for Opponent, if $y \in L_-^\infty$ and $y' \in W_-^\infty$ and $y \subset^- y'$, then there is a necessarily finite and \oplus -maximal configuration $y'' = [\text{Pos}[y]] \in \mathcal{C}^\omega(A)$ such that both $y'' \subset^- y$ and $y'' \subset^- y'$.

Remark Lemma 8 is useful in showing that the class of strategies defined next—concurrent defensive strategies—can be regarded, in a way, as complete for the class of Borel non-zero-sum games for which a Nash equilibrium always exists: every nondeterministic strategy in every strategy profile that forms a Nash equilibrium can be defined with respect to the class of concurrent defensive strategies. It also states that the infinite sets/plays in these games can be regarded as being *finitely decidable* in the following sense: if one can avoid a losing infinite configuration (for Player), then one only needs to avoid some previous finite configuration. Lemma 8 is *essential* for the construction of concurrent winning strategies, which are possible only if nondeterminism is allowed, even in the finite case [13]. From a more non-cooperative viewpoint, winning strategies—and hence Lemma 8—will also be used to construct, so-called, ‘threat’, ‘trigger’, or ‘punishment’ strategies [32, pp. 143], needed to show results about Nash equilibria.

Strategies

Informally, a concurrent defensive strategy always tries to avoid ending in a losing configuration, that is, in a configuration in L_+ (in case it is a strategy σ of Player) or in a configuration in L_- (in case it is a strategy τ of Opponent). Concurrent defensive strategies are defined upon the following, intuitively simple, property: let $\sigma : S \rightarrow A$ be a strategic map of event structures with polarity (i.e., σ respects causal dependency within configurations and injectivity on configurations) such that for all finite configurations $x \in \mathcal{C}^\omega(S)$, with $\sigma x \in L_+$, if there is a configuration $y \in W_+$ such that $\sigma x \subset^+ y$ then there is a configuration $x' \in \mathcal{C}(S)$ such that the following diagram (Diagram †) commutes

$$\begin{array}{ccc} x & \xrightarrow{\quad \subset^+ \quad} & x' \\ \sigma \downarrow & & \downarrow \sigma \\ \sigma x \in L_+ & \xrightarrow{\quad \subset^+ \quad} & y \in W_+ \end{array}$$

and similarly for Opponent.

A concurrent defensive strategy $\sigma : S \rightarrow A$ is a receptive and strategic map of event structures which satisfies Diagram † and for which $(s \in \text{Neg}[x] \ \& \ s' \in (x' \setminus x)) \implies s \rightarrow s'$.

While satisfying Diagram † makes the map defensive, letting the map be strategic and satisfy that $s \rightarrow s'$ as defined above makes the map a strategy, as it is receptive and innocent: immediate causal dependencies in \leq_S not already in \leq_A are from $-ve$ to $+ve$ events. Note, in particular, that since σx in Diagram † is finite then x is finite (as well as $\text{Neg}[x]$). Therefore, no s' is made to causally depend on infinitely many s , and so S is an event structure (in particular, axiom (i) of event structures which requires all events to have a finite set of causes—see Section 2.1—is not violated).

A strategy can also be defined with respect to a subgame A_y of a given game A , with $y \in \mathcal{C}(A)$. Formally, let $G = (A, W_+, W_-)$ be a game. Then for every $y \in \mathcal{C}(A)$ define a *residual subgame* $G_y = (A, W_+, W_-)_y \stackrel{\text{def}}{=} (A_y, W_y^+, W_y^-)$ where:

$$\begin{aligned} A_y &= \{a \in A \setminus y \mid \exists y' \in \mathcal{C}(A). y \subseteq y' \ \& \ a \in y'\} \\ W_y^+ &= \{y' \subseteq A_y \mid y \cup y' \in W_+\} \\ W_y^- &= \{y' \subseteq A_y \mid y \cup y' \in W_-\} \end{aligned}$$

with the order \leq_{A_y} being the restriction of \leq_A to A_y and consistency defined as $\text{Con}_{A_y} = \{Y \subseteq_{\text{fin}} A_y \mid Y \in \text{Con}_A\}$.

Intuitively, a residual subgame (simply a *subgame* hereafter) is the remainder of a concurrent game once a certain position (a configuration) has been reached. Thus, they can be understood as the residual of a game with respect to a particular configuration. Given a game G , we write y -subgame for the subgame G_y of G ; we write σ_y for a strategy, called a y -substrategy, in the y -subgame G_y . Given a configuration z in a y -subgame A_y we also write $z \in W_+$ for $z \in W_+^+$, and similarly for the other winning/losing sets—as well as for sets of subconfigurations.

PROPOSITION 9 (Closure under subgames). For each concurrent game (A, W_+, W_-) , if $\sigma : S \rightarrow A$ is a winning strategy for Player then for every $y \in \mathcal{C}(A)$ such that $y = \sigma x$ for some $x \in \mathcal{C}(S)$, there is a winning strategy σ_y^x in $(A, W_+, W_-)_y$.

To see that Proposition 9 holds over non-zero-sum concurrent games, let $\sigma_y^x : S_x \rightarrow A_y$ be the strategy in the y -subgame A_y defined by the event structure $(S_x, \leq_{S_x}, \text{Con}_{S_x})$, where:

$$S_x \stackrel{\text{def}}{=} \{s \in S \setminus x \mid \exists x' \in \mathcal{C}(S). x \subseteq x' \ \& \ s \in x'\},$$

the order \leq_{S_x} is the restriction of \leq_S to the set S_x , and consistency relation $\text{Con}_{S_x} = \{X \subseteq_{\text{fin}} S_x \mid X \in \text{Con}_S\}$. Thus, the map σ_y^x is the y -substrategy of σ , i.e. a strategy in the y -subgame A_y , with respect to $x \in \mathcal{C}(S)$. Then, for all $x_m \in \mathcal{C}(S_x)$ we have that

$$x_m \text{ is } \ominus\text{-maximal in } S_x \implies x \cup x_m \text{ is } \ominus\text{-maximal in } S.$$

And since for all \ominus -maximal configurations $x' \in \mathcal{C}(S)$ we have that $\sigma x' \in W_+$ then, because of the definition of W_y^+ (i.e. of the y -subgame A_y) we have that $\sigma_y^x x_m \in W_y^+$ for all \ominus -maximal configurations $x_m \in \mathcal{C}(S_x)$. Therefore σ_y^x is a winning strategy for Player in the non-zero-sum game $(A, W_+, W_-)_y$, as desired.

The latter two results, namely, Lemma 8 and Proposition 9, can now be used to construct concurrent defensive strategies in subgames, provided that they are race-free, bounded-concurrent, and determine a Borel set of winning configurations/plays.

LEMMA 10. Let $(A, W_+, W_-)_y$ be a y -subgame of a race-free bounded-concurrent Borel non-zero-sum game (A, W_+, W_-) . If Opponent does not have a strategy τ_y such that $\langle \sigma_y, \tau_y \rangle \subseteq W_-$, for all σ_y of Player, then there is a concurrent defensive strategy σ'_y for Player such that $\langle \sigma'_y, \tau'_y \rangle \subseteq L_-$, for all τ'_y of Opponent.

Proof. (Sketch) As Opponent has no τ_y such that $\langle \sigma_y, \tau_y \rangle \subseteq W_-$, for all strategies σ_y of Player, then Opponent does not have a winning strategy in the $(A, L_-, W_-)_y$. Thus, we know that

$$y \in W_- \implies \exists y' \in L_-. y \subset^+ y',$$

and use Lemma 8 (because now the losing/winning configurations of Opponent are winning/losing for Player) to construct a concurrent defensive strategy σ'_y for which $\langle \sigma'_y, \tau_y \rangle \subseteq L_-$. And, since the game has Borel sets of winning conditions for both players, then we know that the (zero-sum) y -subgame $(A, L_-, W_-)_y$ is determined, i.e. that σ'_y is a winning strategy in $(A, L_-, W_-)_y$ for Player—hence $\langle \sigma'_y, \tau'_y \rangle \subseteq L_-$, for all τ'_y of Opponent. Proposition 9 ensures the existence of winning strategies in subgames. \square

Note that the strategy σ'_y that is constructed in Lemma 10 is not, in general, a winning strategy in $(A, W_+, W_-)_y$ for Player: $\langle \sigma'_y, \tau'_y \rangle \subseteq L_-$, but not necessarily $\langle \sigma'_y, \tau'_y \rangle \subseteq W_+$. However, σ'_y is a winning strategy for Player in the inverse game G^{-1} . We can think of σ'_y as a *punishment* strategy instead, that is, a strategy to ensure that the other player does not achieve its goal. Note also that since Borel zero-sum concurrent games are determined [21], and Lemma 10 shows that whenever there is a winning strategy in a zero-sum concurrent game then there is also a winning concurrent defensive strategy, the following result immediately holds:

COROLLARY 11. *Let G be a race-free bounded-concurrent Borel zero-sum game. Player has a winning strategy in G iff it has a concurrent defensive one.*

Thus far we have made use of winning strategies in various ways; indeed, for zero-sum games the concept of *determinacy* [26] is much more interesting than that of *equilibrium* [32]. However, when players' preferences are binary, the two are in fact strongly related. A corollary of Proposition 3 and Corollary 11 is:

PROPOSITION 12. *Let G be a race-free bounded-concurrent Borel zero-sum game. Then both $NE(G) \neq \emptyset$ and there is a strategy profile $(\sigma, \tau) \in NE(G)$ where either σ or τ is a winning concurrent defensive strategy.*

Games that are non-zero-sum and where $W_+ \cap W_- \neq \emptyset$ also have a Nash equilibrium, in this case, even within the smaller class of deterministic strategies and the larger class of games that are neither race-free nor bounded-concurrent.

PROPOSITION 13. *Let $G = (A, W_+, W_-)$ be a game such that $W_+ \cap W_- \neq \emptyset$. Then $NE(G) \neq \emptyset$. Moreover, there are*

- $(\sigma_0, \tau_0) \in NE(G)$ where σ_0 and τ_0 are deterministic, and
- $(\sigma, \tau) \in NE(G)$ and σ or τ is a concurrent defensive strategy.

Proof. (Sketch) Let $y \in (W_+ \cap W_-)$ and σ and τ be the following inclusion maps:

- $\sigma_0 : A \uparrow (\{a \in A \mid a \in y \text{ or } \text{pol}(a) = -\}) \hookrightarrow A$,
- $\tau_0 : A^\perp \uparrow (\{a \in A^\perp \mid a \in y \text{ or } \text{pol}(a) = +\}) \hookrightarrow A^\perp$.

Then, it follows from their definitions that σ_0 and τ_0 are deterministic, and clearly also that $\{y\} = \langle \sigma_0, \tau_0 \rangle$. By Proposition 3, item 1, we have that $(\sigma_0, \tau_0) \in NE(G)$.

Based on $\sigma_0 : S_0 \rightarrow A$ one can construct a, possibly non-deterministic, concurrent defensive strategy $\sigma : S \rightarrow A$: for each $\sigma_0 x \subseteq y$, if $(\sigma_0 x) \in L_+$ then we re-define S_0 so that the resulting map is a concurrent defensive strategy (by adding extra-causal dependencies from $-ve$ events in x to $+ve$ events in x_y , with $\sigma_0 x_y = y$, as required). Moreover, let $\tau = \tau_0$. Since $\{y\} = \langle \sigma, \tau \rangle$, again, by Proposition 3, we have that $(\sigma, \tau) \in NE(G)$. \square

With respect to the winning sets of Player and Opponent, there is one more case to consider when studying the existence of Nash equilibria of concurrent games on event structures, namely, when $W_+ \cap W_- = \emptyset$ and $L_+ \cap L_- \neq \emptyset$. To show such a case, the following closure property with respect to inverse games is useful:

PROPOSITION 14 (Closure under inverse concurrent games). *Let (A, W_+, W_-) be determined. If there is no winning strategy in (A, W_+, W_-) for Player/Opponent, there is a winning strategy in the inverse game $(A, W_+, W_-)^{-1}$ for Opponent/Player.*

And the main theorem of this section follows:

THEOREM 15 (Nash equilibria). *If G is a race-free and bounded-concurrent Borel non-zero-sum game, then we have that*

$$NE(G) \neq \emptyset.$$

Proof. Let $G = (A, W_+, W_-)$ be a race-free bounded-concurrent Borel non-zero-sum game. With respect to the payoff sets of the game, there are three cases to consider, namely:

1. $W_+ \cap W_- \neq \emptyset$,
2. $W_+ = L_-$,
3. $W_+ \cap W_- = \emptyset$ and $L_+ \cap L_- \neq \emptyset$.

The existence of Nash equilibria in concurrent games on event structures, thus, can be analysed by checking the cases given above: for cases 1 and 2 the result follows from Propositions 13 and 12, respectively. For case 3 we can consider two sub-cases, namely:

- (i) when one of the players has a winning strategy or
- (ii) when neither player has a winning strategy.


Firstly, note that both players cannot have a winning strategy because $W_+ \cap W_- = \emptyset$. Let us now analyse the two cases above. For case (i) the result follows from items 2 and 3 of Proposition 3: if Player has a winning strategy σ in $G = (A, W_+, W_-)$ then we have $\langle \sigma, \tau' \rangle \subseteq W_+$, for all strategies τ' of Opponent. Since $W_+ \cap W_- = \emptyset$ then $W_+ \subseteq L_-$; hence item 2 of Proposition 3 holds. Then every strategy profile (σ, τ') is in $NE(G)$. If, on the other hand, Opponent has a winning strategy τ in $G = (A, W_+, W_-)$ then we have $\langle \sigma', \tau \rangle \subseteq W_-$, for all strategies σ' of Player. Again, since $W_+ \cap W_- = \emptyset$ then $W_- \subseteq L_+$; hence item 3 of Proposition 3 holds. Then every strategy profile (σ', τ) is in $NE(G)$.

For case (ii) the result follows from item 4 of Proposition 3 by letting both Player and Opponent use winning strategies in the inverse game G^{-1} : If no player has a winning strategy in G then, by Proposition 14, both players have winning strategies in the inverse game $G^{-1} = (A, L_-, L_+)$. Let σ and τ be such winning strategies. Then, it follows that $\langle \sigma, \tau \rangle \subseteq (L_+ \cap L_-)$. Moreover, since σ (resp. τ) is a winning strategy in G^{-1} then $\langle \sigma, \tau' \rangle \subseteq L_-$ (resp. $\langle \sigma', \tau \rangle \subseteq L_+$), for all σ' of Player. Hence item 4 of Proposition 3 holds. Then, the strategy profile (σ, τ) is in $NE(G)$, as desired. \square

From Theorem 15 it follows that the class of concurrent defensive strategies is, in a sense, complete for the existence of Nash equilibria in Borel non-zero-sum concurrent games on event structures. Note that to prove cases 1–3 in Theorem 15, in each case a concurrent defensive strategy is constructed.

COROLLARY 16. *Let G be a race-free bounded-concurrent Borel non-zero-sum game. There is $(\sigma, \tau) \in NE(G)$, where σ or τ is a concurrent defensive strategy.*

EXAMPLE 17. *Let $G = (A, W_+, W_-)$ be the concurrent game given below, a system similar in spirit to that in Figure 1. In G we have the following winning sets: $W_+ = \{y \in \mathcal{C}(A) \mid \oplus_B \in y\}$ and $W_- = \{y \in \mathcal{C}(A) \mid \oplus_R \in y \Leftrightarrow \ominus_A \in y\} \setminus \emptyset$. Note that the game is non-zero-sum and no player has a winning strategy. However, due to Theorem 15 there is a Nash equilibrium—and because of Corollary 16 we know that at least one with a concurrent defensive strategy—where, in this case, both players achieve their goals:*



It should be easy to see that the only result of play (for $\sigma \otimes \tau$) in the game G is $\oplus_R \rightarrow \ominus_A \rightarrow \oplus_B$, the desired behaviour of the system we wish to synthesize. In other words, the (internal) event structures S and T of σ and τ are the desired “computation plans” for the system and the environment, respectively. \square

Remark One may be wondering why if both race-freedom and bounded-concurrency already hold in the best-known concurrent games in the computer science literature, e.g. those games in [8], the existence of Nash equilibrium strategy profiles does not hold in nondeterministic concurrent strategies, as in our case. The reason, as mentioned before, is that such concurrent games are assumed to be *synchronous* in the following sense: all players are assumed to make simultaneous choices in each round (like in the matching pennies game), with no power to asynchronously wait—possibly forever—until its opponent plays. Indeed, *round-free asynchrony* (not studied before) as defined in our model of concurrent games makes the Nash equilibrium problem very different.

6. Dynamic Equilibria

As shown in Example 4, sometimes there are “undesirable” Nash equilibria, say y , in a game G : results y such that both $y' >_+ y$ and $y' >_- y$ for some $y' \in NE(G)$; in Example 4, for instance, $\{\oplus_R, \ominus_R\} >_+ \{\oplus_L, \ominus_L\}$ and $\{\oplus_R, \ominus_R\} >_- \{\oplus_L, \ominus_L\}$, *i.e.*, when playing (σ_L, τ_L) both players get payoffs that are worse than what they would have obtained had they played (σ_R, τ_R) instead.

This is a well-known anomaly of Nash equilibria when playing games with multiple rounds. Such games are sometimes called ‘repeated’ or ‘dynamic’ games in the game theory literature [32]. A refinement of Nash equilibria that deals with this problem is *subgame perfect Nash equilibria*. We say that a Nash equilibrium is subgame perfect iff it is a Nash equilibrium in all subgames.

For instance, in Example 4, the strategy profile (σ_L, τ_L) is not subgame perfect because $((\sigma_L)_{\{\oplus_R\}}, (\tau_L)_{\{\oplus_R\}})$ is not a Nash equilibrium in the subgame $A_{\{\oplus_R\}}$. We say that σ_y is a y -substrategy of σ , that is, σ restricted to the y -subgame A_y .

Since we have a well-defined notion of subgame it is only natural to ask whether every game on an event structure has a subgame perfect Nash equilibrium. We answer this question positively. Let $SPNE(G)$ be the set of subgame perfect Nash equilibria of the game G . Since $SPNE(G)$ is a refinement of $NE(G)$ we have that $SPNE(G) \subseteq NE(G)$. The following observation is important.

A player has an incentive to deviate from an initially chosen strategy if a better payoff can be achieved by making different choices once having reached a particular subgame. Because of the binary nature of players’ preferences, Player/Opponent has an incentive to deviate only if a result in W_+/W_- can be achieved, and has no incentive to deviate if its payoff is already in W_+/W_- . An analysis of the winning sets leads to the following result.

THEOREM 18 (Subgame perfect Nash equilibria). *If G is a race-free and bounded-concurrent Borel non-zero-sum game, then*

$$SPNE(G) \neq \emptyset.$$

Proof. (Sketch) As for Theorem 15, let $G = (A, W_+, W_-)$ and consider the same three cases before, namely: (1) $W_+ \cap W_- \neq \emptyset$, (2) $W_+ = L_-$, or (3) $W_+ \cap W_- = \emptyset$ and $L_+ \cap L_- \neq \emptyset$.

In case 1 neither player unilaterally deviates because their payoffs are already winning—so best responses in all subgames (Nash equilibria in all subgames). In case 2, again, we construct a winning strategy for one of the players, who will not deviate because its payoff is winning. The other player cannot improve its payoff, which is losing. Any such strategy profile is a Nash equilibrium in all subgames as one of the players is using a winning strategy—hence also winning in all subgames. In case 3 an analysis similar to that in Theorem 15 applies: either one of the players has a winning strategy (so case 2 analysis applies) or let both players use winning strategies in the inverse game G^{-1} , which are winning strategies in all subgames of G^{-1} —hence Nash equilibria in all subgames. \square

Since in Theorem 18 we make use of a deterministic or a winning strategy, then the next result immediately follows:

COROLLARY 19. *Let G be a race-free bounded-concurrent Borel non-zero-sum game. There is $(\sigma, \tau) \in SPNE(G)$ where σ or τ is a concurrent defensive strategy.*

Remark It is worth noting that, contrarily to the computer science literature about equilibria on *interleaving* models of concurrency (such as transition systems or concurrent game structures), the literature about *noninterleaving* models is very scarce. In fact, results such as Theorems 15 and 18 ensuring the existence of equilibria within other noninterleaving models, *e.g.* Petri nets, asynchronous transition systems or Mazurkiewicz traces, are not known (unless their interleaving semantics are considered).

7. Examples and Applications


As mentioned before, our two structural conditions on games, namely, race-freedom and bounded-concurrency, mean no actual restriction with respect to usual applications in semantics and logic (*e.g.*, model-checking, satisfiability, etc.) as well as best-known game models of logics and (concurrent) systems with interleaving semantics. To be more precise a few examples are given next.

Semantics of fixpoint logics First, observe that our main focus on race-freedom is not overly restrictive in logical applications. For example, a translation of parity games to games on event structures delivers the standard parity game semantics of fixpoint logics such as the μ -calculus [25]. The translation constructs a tree-like infinite event structure where all branches are in conflict and each finite configuration is either \oplus -maximal (and in $W_+ \cap L_-$) or \ominus -maximal (and in $W_- \cap L_+$), thus implicitly introducing rounds to the game and forcing both players to play to maximal and infinite configurations in order to win. The event structure above is both race-free and bounded-concurrent and the game is zero-sum.

Semantics of logics for multi-agent systems The translation above can be adapted to represent (give semantics of) concurrent game structures with complete information, which are used for the denotation of alternating-time μ -calculus (AMC) formulae [3]. In this case, the “protagonist” of the AMC game is modelled by Player and the “antagonist” is modelled by Opponent. Winning conditions and rounds are introduced as before, wherein the moves of Player and Opponent strictly alternate and with Player always moving/playing first in this verification game.

Note that the two examples above do not exploit at all the non-zero-sum power of our model of games. They are presented here only to illustrate that both race-freedom and bounded-concurrency do not represent a semantic limitation to capture the behaviour of other important and complex game models in the literature. The next examples show, on the other hand, situations where having non-zero-sum expressive power in the game can be useful.

Reasoning about strategic and cooperative behaviour In 1959 Robert Aumann showed that sometimes cooperation can be strategically/rationally achieved/sustained only when infinite behaviour is considered—that is when players can meet infinitely often [4]. *Non-zero-sum games are required to model cooperation.* Consider the “Peace-Attack” game where Player and Opponent can choose to either attack or make peace with the other. In such a game a player loses (i) if no proposal is made, (ii) if peace is proposed when the other proposes to attack—in which case the other wins, (iii) if peace is made only a finite number of times, or (iv) if both propose to attack always. They both can win only if peace is proposed infinitely often—*i.e.*, always. Formally the game is as illustrated below:



$W_+ = \{(\oplus_P, \ominus_P)^\infty, (\oplus_P, \ominus_P)^* \oplus_A, (\oplus_A, \oplus_A)^* (\oplus_A, \ominus_P)\}$, and $W_- = \{(\oplus_P, \ominus_P)^\infty, (\oplus_A, \oplus_A)^* \oplus_P, (\oplus_P, \ominus_P)^* (\oplus_P, \oplus_A)\}$.

It can be seen that no player in this non-zero-sum game can unilaterally ensure to end in a winning configuration. It can also be

shown that *always attacking*—which ends in a losing configuration for both players—is a Nash equilibrium of any *finite* version of the game, which must only contain plays of even length. Yet, if we assume that players are acting selfishly, but rationally, then cooperation can be sustained infinitely often by the two players *always proposing peace*, the desired behaviour (and the *only* Pareto efficient [32] subgame perfect Nash equilibrium) of the system! This kind of strategic reasoning, namely about infinite and non-zero-sum games, is formally captured by the (Nash) Folk theorems in the classical game theory literature [32].

Semantics of logics for concurrency Our non-zero-sum games provide a model for Sandu and Hintikka’s Independence-Friendly (IF) logic [23]. IF logic is an extension of classical first-order (FO) logic where quantifiers are partially rather than totally ordered.

The syntax and semantics of IF logic is as follows. The syntax of IF logic is that of classical first-order logic with quantifiers extended to *slashed quantifiers* $\exists x/y, z, \dots, \forall x/y, z, \dots$, whose intended interpretation is that when choosing the witness/counterexample x , we may not know the values of y, z, \dots (which are presumed to be bound earlier in the formula). For example, in

$$\forall x. \exists y. \forall u/x, y. \exists v/x, y : \varphi(x, y, u, v)$$

the variable v depends only on u , not on x or y (and y , of course, depends only x because u, v have not yet been mentioned). In other words, the semantics is given by a game in a partially ordered structure where $x \leq y$ and $u \leq v$. In general, these dependencies define a partial order which, under our concurrent game semantics, is straightforwardly interpreted with an event structure with polarity.

We should note that our games are not quite IF, as we require dependency to be transitive. The standard definition of IF allows for non-transitive dependencies between quantifiers, which does not increase the expressive power of the logic, but does result in rather pathological properties of formulae. For this reason, some IF work imposes transitivity as an additional requirement.

In our IF variant, called the *lazy IF logic*, neither player—as in our event structure games framework—is forced to play; the arena game A is, as usual, defined by the formula φ at hand using the explicit variable dependencies in the IF formula (*cf.*, see [14] too), and with winning conditions as follows: let $y \in \mathcal{C}(A)$; we define

1. $y \in L_+$ & $y \in L_-$, if $\overline{max}_+(y)$ & $\overline{max}_-(y)$;
2. $y \in L_+$ & $y \in W_-$, if $\overline{max}_+(y)$ & $max_-(y)$;
3. $y \in W_+$ & $y \in L_-$, if $max_+(y)$ & $\overline{max}_-(y)$;
4. $y \in W_+$ & $y \in L_-$, if $max_+(y)$ & $max_-(y)$ & $\varphi(y)$ true;
5. $y \in L_+$ & $y \in W_-$, if $max_+(y)$ & $max_-(y)$ & $\varphi(y)$ false;

where $\varphi(y)$ is formula φ evaluated using the unique variable values/assignment given by configuration y . As y is both \oplus -maximal and \ominus -maximal, it is necessarily a maximal configuration of A . A model M of a lazy-IF formula φ is a non-empty and finite universe of values V_M . Thus, the truth of formula φ is defined with respect to an assignment of values in V_M . Let $G_{(M, \varphi)}$ be the concurrent game given by the lazy-IF formula φ with respect to a model M . The event structure game semantics of a lazy-IF formula φ , written as $[[\varphi]]_M$, is given by any $(\sigma, \tau) \in SPNE(G_{(M, \varphi)})$.

It is known [14] that IF logic formulae determine event structures that are race-free; and since we only consider finite universes of values V_M , then such event structures contain only finite configurations (hence are bounded-concurrent). Then, by Theorem 18, such games *always* have a subgame perfect Nash equilibrium.

These results justify/allow the following definition: let φ be a lazy-IF formula and M a model. Then, for all Pareto efficient $(\sigma, \tau) \in SPNE(G_{(M, \varphi)})$, and all $y \in \langle \sigma, \tau \rangle$:

1. φ is true iff $y \in W_+$ & $y \in L_-$;
2. φ is false iff $y \in L_+$ & $y \in W_-$;
3. φ is undetermined iff $y \in L_+$ & $y \in L_-$.

EXAMPLE 20. Let φ be the lazy-IF formula $\forall x. \exists y/x : x = y$ and M be given by $dom(x) = dom(y) = \{0, 1\}$. This formula states that for all values of x there is a value of y that is independent of x such that $x = y$ holds. Seen as an IF formula, φ is undetermined. Assuming that the goal of Player is to show that φ is true and that the goal of Opponent is to show that φ is false, we can infer that no player has a winning strategy to satisfy its goal. However $G_{(M, \varphi)}$ has a subgame perfect Nash equilibrium (σ, τ) , namely the one where each player waits for the other to play first. In such a case the only outcome $y \in \langle \sigma, \tau \rangle$ of the game is $y = \emptyset$, which is in $L_+ \cap L_-$ according to the definition of the game, and characterising φ as undetermined (see also Proposition 21 below). Thus, A is

$$\oplus_0 \rightsquigarrow \oplus_1 \qquad \ominus_0 \rightsquigarrow \ominus_1$$

with $W_+ = \{\{\oplus_0\}, \{\oplus_1\}, \{\oplus_0, \oplus_1\}, \{\oplus_1, \oplus_1\}\}$ and, similarly, with $W_- = \{\{\ominus_0\}, \{\ominus_1\}, \{\ominus_0, \ominus_1\}, \{\oplus_1, \ominus_0\}\}$. \square

IF logic and the lazy-IF logic presented in this section have the same syntax but a “different” game semantics. However:

PROPOSITION 21. For all IF formulae φ , we have φ true iff it is lazy-IF true; φ false iff it is lazy-IF false; and φ undetermined iff it is lazy-IF undetermined.

8. Analysis and related work

From a semantic point of view, models for concurrency have traditionally been classified as either interleaving or noninterleaving. This paper is, to the best of our knowledge, the first one to provide positive results about the existence of equilibria for a noninterleaving model of concurrency, namely for event structures. Our model of concurrent games is round-free, nondeterministic, and asynchronous (interpreted in a noninterleaving way)—a semantically novel combination of features with respect to other game models in the computer science, AI, and game theory literatures.² Our results start to complement the great many results already known for other definitions of concurrency and independence in games, *e.g.* those on graphs, trees, or concurrent game structures (*cf.*, [8, 27]).

The main results of this paper (which generalise the main results in [13, 21, 22]) characterise the existence of equilibria in games on event structures. They reveal that race-freedom and bounded-concurrency are necessary and sufficient. As these two conditions implicitly hold in previously studied (interleaving) game models of concurrency, it was not known that they were actually needed for the existence of equilibria. Our results establish the boundaries for the amount of ‘true-concurrency’ that can be allowed when analysing problems about equilibria. Specifically, our results apply to non-zero-sum *concurrent* games with perfect information over a noninterleaving model of concurrency. For a good survey of concurrent games on interleaving models or of (stochastic) games with imperfect information we refer the reader to [8].

We should note that having binary payoff sets and nondeterministic strategies—instead of mixed ones (*i.e.*, randomized)—may be positive. It is known that randomization, either in the strategies [37] or the arenas [36] can lead to undecidability; however, w.r.t. binary payoff sets, decidability can be recovered [35]. Although we did not study decidability issues—let alone complexity questions—this restriction is still very natural, especially when pure-strategy equilibria are considered [6, 35] and applications to verification are to be considered (*e.g.*, non-zero-sum games on Petri nets). Positive results should also be achieved for well-founded games [13] with quantitative payoffs [12] and probabilistic strategies [42].

²Many other different definitions of independence and concurrency in games (*e.g.*, asynchronous, distributed, etc.) can be found in the literature; we reference only those models of games for which problems about Nash equilibrium and subgame perfect Nash equilibrium have been studied.

From a more classical game theoretic perspective, two topics deserve to be discussed: multi-player and stochastic games. Firstly, our focus on two-player games only can simplify some computational situations; for instance, in certain scenarios (e.g., two-player non-zero-sum infinite games on graphs), the existence of Nash equilibria is known for two-player games, but such results do not immediately transfer to three-player games (cf., [7, 8, 38]). Also, as mentioned before, the use of stochastic strategies does not ensure the existence of Nash equilibria in concurrent games on event structures that are not race-free or bounded-concurrent. However, with respect to other definitions of concurrency or arena games, having the power of stochastic behaviour can be quite beneficial; for instance, it can be used to ensure the existence of equilibria, even for some multi-player games (e.g., [8, 9]). All these results show that even though our focus on two-player, perfect-information, non-zero-sum concurrent Borel games with pure strategies is natural, further investigations into more general games are still needed.

Future work It is often the case that results on event structures can be imported to other models of concurrency. It would be interesting to see whether some of our results can be transferred to other noninterleaving models of concurrency—such as Petri nets, asynchronous transition systems, or Mazurkiewicz traces—via the adjunctions with those formalisms and the more mathematically abstract presentations of concurrent strategies as profunctors [41].

Acknowledgments

We thank Julian Bradfield, Paul Harrenstein, Glynn Winskel, and the reviewers for helpful comments as well as the ERC Advanced Research Grant “RACE” (291528) at Oxford for financial support.

References

- [1] S. Abramsky. Semantics of interaction: an introduction to game semantics. In *Newton Institute Publications*. Camb. Univ. Press, 1996.
- [2] S. Abramsky. Sequentiality vs. concurrency in games and logic. *Mathematical Structures in Computer Science*, 13(4):531–565, 2003.
- [3] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
- [4] R. Aumann. Acceptable points in general cooperative n-person games. In *Contributions to the Theory of Games IV*, volume 40 of *Annals of Mathematics Study*, pages 287–324. Princeton Univ. Press, 1959.
- [5] J. v. Benthem. Logic games: From tools to models of interaction. In *Proof, Computation and Agency*, volume 352 of *Synthese library*, pages 183–216. Springer, 2011.
- [6] P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Concurrent games with ordered objectives. In *FoSSaCS*, volume 7213 of *LNCS*, pages 301–315. Springer, 2012.
- [7] K. Chatterjee. Two-player nonzero-sum ω -regular games. In *CONCUR*, volume 3653 of *LNCS*, pages 413–427. Springer, 2005.
- [8] K. Chatterjee and T. A. Henzinger. A survey of stochastic ω -regular games. *Journal of Computer and System Sciences*, 78(2):394–413, 2012.
- [9] K. Chatterjee, R. Majumdar, and M. Jurdzinski. On Nash equilibria in stochastic games. In *CSL*, volume 3210 of *LNCS*, pages 26–40. Springer, 2004.
- [10] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theoretical Computer Science*, 365(1-2):67–82, 2006.
- [11] K. Chatterjee, T. A. Henzinger, and B. Jobstmann. Environment assumptions for synthesis. In *CONCUR*, volume 5201 of *LNCS*, pages 147–161. Springer, 2008.
- [12] P. Clairambault and G. Winskel. On concurrent games with payoff. *ENTCS*, 83:61–84, 2013.
- [13] P. Clairambault, J. Gutierrez, and G. Winskel. The winning ways of concurrent games. In *LICS*, pages 235–244. IEEE, 2012.
- [14] P. Clairambault, J. Gutierrez, and G. Winskel. Imperfect information in logic and concurrent games. In *Computation, Logic, Games, and Quantum Foundations*, volume 7860 of *LNCS*, pages 7–20. Springer, 2013.
- [15] J. Esparza and K. Heljanko. *Unfoldings - A Partial-Order Approach to Model Checking*. Monographs in TCS. Springer, 2008.
- [16] D. R. Ghica. Applications of game semantics: From program analysis to hardware synthesis. In *LICS*, pages 17–26. IEEE, 2009.
- [17] P. Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems*, volume 1032 of *LNCS*. Springer, 1996.
- [18] E. Grädel. Model checking games. *ENTCS*, 67:15–34, 2002.
- [19] E. Grädel and M. Ummels. Solution concepts and algorithms for infinite multiplayer games. In *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 151–178, 2008.
- [20] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games*, volume 2500 of *LNCS*, 2002. Springer.
- [21] J. Gutierrez and G. Winskel. Borel determinacy of concurrent games. In *CONCUR*, volume 8052 of *LNCS*, pages 516–530. Springer, 2013.
- [22] J. Gutierrez and G. Winskel. On the determinacy of concurrent games on event structures with infinite winning sets. *Journal of Computer and System Sciences*, DOI: 10.1016/j.jcss.2014.04.005, 2014.
- [23] J. Hintikka and G. Sandu. A revolution in logic? *Nordic Journal of Philosophical Logic*, 1(2):169–183, 1996.
- [24] M. Hyland. Game semantics. In *Newton Institute Publications*. Camb. Univ. Press, 1997.
- [25] D. Kozen. Results on the propositional mu-calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [26] D. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
- [27] D. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- [28] J. Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, 1951.
- [29] M. Nielsen and G. Winskel. Models for concurrency. In *Handbook of Logic in Computer Science*, pages 1–148. Oxf. Univ. Press, 1995.
- [30] M. Nielsen, G. D. Plotkin, and G. Winskel. Petri nets, event structures and domains, Part I. *Theoretical Computer Science*, 13:85–108, 1981.
- [31] N. Nisan. Introduction to mechanism design. In *Algorithmic Game Theory*, pages 209–242. Camb. Univ. Press, 2007.
- [32] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.
- [33] S. Rideau and G. Winskel. Concurrent strategies. In *LICS*, pages 409–418. IEEE, 2011.
- [34] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Camb. Univ. Press, 2008.
- [35] M. Ummels and D. Wojtczak. Decision problems for Nash equilibria in stochastic games. In *CSL*, volume 5771 of *LNCS*, pages 515–529. Springer, 2009.
- [36] M. Ummels and D. Wojtczak. The complexity of Nash equilibria in simple stochastic multiplayer games. In *ICALP*, volume 5556 of *LNCS*, pages 297–308. Springer, 2009.
- [37] M. Ummels and D. Wojtczak. The complexity of Nash equilibria in limit-average games. In *CONCUR*, volume 6901 of *LNCS*, pages 482–496. Springer, 2011.
- [38] N. Vieille. Two player stochastic games II: the case of recursive games. *Israel Journal of Mathematics*, 119:93–126, 2000.
- [39] I. Walukiewicz. A landscape with games in the background. In *LICS*, pages 356–366. IEEE, 2004.
- [40] G. Winskel. Event structures. In *Advances in Petri Nets*, volume 255 of *LNCS*, pages 325–392. Springer, 1986.
- [41] G. Winskel. Strategies as profunctors. In *FoSSaCS*, volume 7794 of *LNCS*, pages 418–433. Springer, 2013.
- [42] G. Winskel. Distributed probabilistic strategies. *ENTCS*, 83:379–392, 2013.