

Observably Deterministic Concurrent Strategies and Intensional Full Abstraction for Parallel-or*

Simon Castellan¹, Pierre Clairambault¹, and Glynn Winskel²

- 1 Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP
first.last@ens-lyon.fr
- 2 Computer Laboratory, University of Cambridge
Glynn.Winskel@cl.cam.ac.uk

Abstract

Although Plotkin’s *parallel-or* is inherently deterministic, it has a non-deterministic interpretation in games based on (prime) event structures – in which an event has a unique causal history – because they do not directly support disjunctive causality. *General event structures* can express disjunctive causality and have a more permissive notion of determinism, but do not support hiding. We show that (structures equivalent to) *deterministic* general event structures do support hiding, and construct a new category of games based on them with a deterministic interpretation of $\mathbf{aPCF}_{\text{por}}$, an affine variant of \mathbf{PCF} extended with parallel-or. We then exploit this deterministic interpretation to give a relaxed notion of determinism (observable determinism) on the plain event structures model. Putting this together with our previously introduced concurrent notions of well-bracketing and innocence, we obtain an intensionally fully abstract model of $\mathbf{aPCF}_{\text{por}}$.

1998 ACM Subject Classification F.3.2 Semantics of Programming Languages

Keywords and phrases Game semantics, parallel-or, concurrent games, event structures, full abstraction

Digital Object Identifier 10.4230/LIPIcs.FSCD.2017.12

1 Introduction

Plotkin’s *parallel-or* is, in a sense, the most well-understood of programming language primitives. Recall that it is a primitive $\text{por} : \mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B}$ defined by the three equations:

$$\text{por } tt \perp = tt \qquad \text{por } \perp t = tt \qquad \text{por } ff = ff$$

As Plotkin famously proved [16], it is the primitive to be added to the paradigmatic purely functional higher-order programming language \mathbf{PCF} in order to get a perfect match (*full abstraction*) with respect to Scott domains. Since Plotkin’s result, full abstraction has been the gold standard for semanticists. Indeed, whereas an adequate model is always *sound* to reason about program equivalence; a *fully abstract model* is also *complete*: two programs are observationally equivalent if and only if they have the same denotation. Though since Plotkin’s paper fully abstract models of various programming languages have been proposed (often through game semantics, with *e.g.* state [3, 1], control [5, 14], concurrency [11], ...); they are usually quite a bit more complicated than plain Scott domains. Parallel-or is well-understood, in that it has a simple input-output (*extensional*) behaviour, and its presence in \mathbf{PCF} reduces observational equivalence to input-output behaviour.

* This work was partially supported by French LABEX MILYON (ANR-10-LABX-0070), and by the ERC Advanced Grant ECSYM.



But is it *really* so well-understood? While parallel-or is simple *extensionally*, how is it best understood *intensionally*? What is the operational behaviour of $\text{por } M N$? Any programmer will sense the subtleties in implementing such a primitive. It immediately spawns two threads, starting evaluating M and N *in parallel*. If both terminate with ff , it terminates with ff as well after both have finished. But it suffices that *one* of the threads returns tt , for parallel-or to return tt immediately, discarding the other. If *both* return tt then one of the threads will be discarded, but which one depends on the scheduler. From an input-output perspective it does not matter which one is selected as they race to produce the same result, but the race nonetheless happens from an operational viewpoint; and a sufficiently intensional semantics will show it. Such behaviour is useful in practice, for instance to speed up deciding the existence of a “good” branch in a search tree by spawning a thread for each branch to explore.

From a game semantics perspective, understanding this combination of racy concurrency and deterministic extensional behaviour has been an unexpected challenge. In earlier work, we dealt with deterministic parallelism using games and strategies based on event structures [17], showing an intensional full abstraction result for a parallel implementation of **PCF** [8]. We also showed how our tools supported shared memory concurrency [6]. Given that, it is no surprise that the racy behaviour mentioned above can be represented with event structures, yielding an adequate model of *e.g.* **PCF** plus parallel-or. But game semantics based on event structures is *causal*: each event comes with a unique causal history. Dependency is *conjunctive*: an event can only occur after *all* its dependencies. This feature is key for the notion of concurrent innocence leading to our definability result [8] and in fact for the very construction of our basic category of games [7]. This has the unfortunate consequence of forcing the strategy for parallel-or to be *non-deterministic*: the race in the operational behaviour of $\text{por } \text{tt } \text{tt}$ yields a choice between two events competing to return tt , each depending on one argument of por . But interpreting por through non-determinism really is a workaround, as appears through the resulting failure of full abstraction.

Constructing a deterministic intensional model for parallel-or has proved demanding. Plays-based models [11] fail as they inline the non-deterministic choice of the scheduler, and cannot express even pure parallelism deterministically. We saw above that plain event structure games do not work either. Giving a deterministic model for parallel-or involves modifying the latter to allow *disjunctive causality*: a given event may have several distinct causal histories, and an event occurrence does not carry information on its specific causal history. Alternatives to event structures allowing this have existed for a long time: *general event structures* [19]. However, we will see that they do not, in general, support *hiding* – which is required to build a category of strategies. We will show that they do support it modulo some further conditions, making them adequate to model parallel-or deterministically.

Beyond the historical twist of giving a fully abstract games model for a language with parallel-or (as game semantics was originally driven by the full abstraction problem for **PCF** *without* parallel-or [13, 2]), a treatment of disjunctive causality is indispensable in a complete game semantics framework for concurrency. The following example illustrates how mundane and widespread it is: the causal history of a packet arriving from the network is simply its *route*. Clearly, introducing an event for each route is best avoided if possible – especially since all further events depending on it will be hereditarily duplicated as well. Deterministic models for disjunctive causality serve two purposes: they allow for (1) a more general, less intensional notion of determinism, and (2) a coarser equivalence relation between strategies that abstracts away benign races. These issues, that we address in this paper, have become recurrent obstacles in our research programme; and the historical importance of parallel-or in semantics made it the perfect candidate to test and showcase the solution.

Contributions. Concretely, we build an intensionally fully abstract games model of $\mathbf{aPCF}_{\text{por}}$, an *affine* version of \mathbf{PCF} with parallel-or. Our tools are designed with the extension in the presence of *symmetry* [8] to full $\mathbf{PCF}_{\text{por}}$ in mind; but presenting them in an affine case allows us to focus on the issues pertaining to parallel-or and disjunctive causality, orthogonal to symmetry. Presenting everything at once in conference format is not reasonable.

Full abstraction for $\mathbf{aPCF}_{\text{por}}$ has two facets: on the one hand, we need to import the conditions of innocence and well-bracketing from [8], which rely on *conjunctive* causality. On the other hand, we need a disjunctive notion of determinism. Consequently our model will involve: (1) the standard category \mathbf{CG} of concurrent games on event structures, and (2) the main novelty of our paper, a new category \mathbf{Disj} supporting disjunctive determinism. Glueing the two together, we can import in \mathbf{CG} the notion of determinism from \mathbf{Disj} , thereafter dubbed *observational determinism*, and prove intensional full abstraction.

Related work. *Game semantics* is a branch of denotational semantics. Originally driven by the full abstraction problem for \mathbf{PCF} , it has grown in the past 25 years into a powerful and versatile methodology to construct compositionally intensional representations of program execution. This lead on the one hand to many full abstraction results (particularly striking in the presence of state as the fully abstract models are then effectively presentable [3, 1]), and on the other hand to applications, ranging from program analysis (model-checking or equivalence checking) to compilation and hardware synthesis [10].

Within game semantics, the present paper is part of a line of work on so-called “truly concurrent” game semantics first pushed by Melliès and colleagues [4, 15], and which advocates the use of causal structures such as event structures and asynchronous transition systems in interactive semantics of programming languages. This line of work has seen a lot of activity in the past five years, prompted in part by a new category of games and strategies on event structures generalizing all prior work, introduced by Rideau and the last author [17]; but also the presheaf-based framework for concurrent games by Hirschowitz and colleagues [12].

Finally, our account of disjunctive causality relies heavily on determinism for *hiding* to work. The third author and Marc de Visme have developed *event structures with disjunctive causality (edc)* [9], supporting both disjunctive causality and hiding in a non-deterministic setting. Links between these two approaches are being explored; at the very least, extra axioms would have to be imposed on edcs in order to mimic the work here.

Outline. In Section 2 we introduce $\mathbf{aPCF}_{\text{por}}$ and give its (non-deterministic) event structure games model \mathbf{CG} . In Section 3 we introduce the deterministic model \mathbf{Disj} for parallel-or. Finally in Section 4 we glue the two, generalize the conditions of [8] and prove full abstraction.


2 Causal game semantics for affine PCF with parallel-or

2.1 The language $\mathbf{aPCF}_{\text{por}}$

To alleviate notation, we only take booleans \mathbb{B} as base type. We start with **affine PCF** (\mathbf{aPCF}). Its types are either \mathbb{B} , or $A \multimap B$ for some types A and B . Its terms are those of the affine λ -calculus with a divergence \perp and boolean primitives (constants and conditionals):

$$M, N ::= \lambda x. M \mid M N \mid x \mid \text{tt} \mid \text{ff} \mid \text{if } M \text{ then } N_1 \text{ else } N_2 \mid \perp$$

We skip the affine typing rules, which are standard – the typing of $\text{if } M \text{ then } N_1 \text{ else } N_2$ is additive, *i.e.* N_1 and N_2 may share resources, and the N_i have boolean type. The



■ **Figure 1** Two strategies on $\mathbb{B}_1 \multimap \mathbb{B}_2 \multimap \mathbb{B}_3$, for parallel-lor and parallel-or

(call-by-name) operational semantics, also standard, yield an evaluation relation $M \Downarrow v$ between closed terms and *values* (booleans or abstractions). We write $M \Downarrow$ when there is some v such that $M \Downarrow v$, and $M \Uparrow$ otherwise. Two terms $\Gamma \vdash M, N : A$ are **observationally equivalent** ($M \simeq_{\text{obs}} N$) iff for all contexts $\mathcal{C}[-]$ such that $\mathcal{C}[M], \mathcal{C}[N]$ are closed terms of type \mathbb{B} , $\mathcal{C}[M] \Downarrow \Leftrightarrow \mathcal{C}[N] \Downarrow$. Recall that an interpretation $\llbracket - \rrbracket$ in some model \mathcal{M} is **fully abstract** if for all M, N , $M \simeq_{\text{obs}} N$ iff $\llbracket M \rrbracket = \llbracket N \rrbracket$ – it is **intensionally fully abstract** if \mathcal{M} quotiented by the semantic equivalent of \simeq_{obs} is fully abstract.

In [16], Plotkin proved that Scott domains are fully abstract not for **PCF**, but for its extension with the so-called **parallel-or** operation:

$$\frac{M \Downarrow \text{tt}}{\text{por } M \ N \Downarrow \text{tt}} \quad \frac{N \Downarrow \text{tt}}{\text{por } M \ N \Downarrow \text{tt}} \quad \frac{M \Downarrow \text{ff} \quad N \Downarrow \text{ff}}{\text{por } M \ N \Downarrow \text{ff}}$$

The combinator **por** is *not sequential*: it is not the case that it evaluates one of its arguments first. In particular, $\text{por } \text{tt} \perp \Downarrow \text{tt}$ and $\text{por } \perp \text{tt} \Downarrow \text{tt}$.


2.2 Strategies as event structures

As usual in game semantics, computation is a dialogue between the program and its environment. The *moves*, or *events*, are either due to the program (Player, +) or its environment (Opponent, -). They are either variable calls (Questions, \mathcal{Q}) or returns (Answers, \mathcal{A}). Unlike traditional game semantics, in our line of work such dialogues are *partially ordered*.

Figure 1 presents two concurrent strategies, both playing on (the game for) $\mathbb{B}_1 \multimap \mathbb{B}_2 \multimap \mathbb{B}_3$, where subscripts are for disambiguation. The diagrams are read from top to bottom. The first event in both strategies is an Opponent question on \mathbb{B}_3 , initiating the computation. Then Player (the program) starts evaluating in parallel its two arguments (q_1 and q_2). These may return **tt** or **ff**, the wiggly line indicating that they cannot reply *both*. Depending on these, Player may eventually answer q_3 . In both diagrams, ff_3 requires *both* arguments to evaluate to **ff**; however they differ as to the events that trigger an answer tt_3 . These diagrams will be made more formal later, but we invite the readers to examine them and convince themselves that the first diagram represents a parallel implementation of the *left or* (diverging if its first argument diverges), whereas the second diagram represents *parallel or*.

Event structures. Such diagrams are formalized as *event structures*. We use here event structures with binary conflict, whereas those of [17, 7] have a more general set of *consistent sets*. Binary conflict is sufficient for our purposes, and preserved by all operations we need.

► **Definition 1.** A (prime) **event structure** (with binary conflict, **es** for short) is $(E, \leq_E, \#_E)$ where E is a set of events, \leq_E is a partial order on E called **causality** and $\#_E$ is an irreflexive symmetric binary relation called *conflict*, such that:



■ **Figure 2** Representations of the arenas $\llbracket \mathbb{B} \rrbracket$ and $\llbracket \mathbb{B}_1 \multimap \mathbb{B}_2 \multimap \mathbb{B}_3 \rrbracket$.

- $\forall e \in E$, $[e] = \{e' \in E \mid e' \leq_E e\}$ is finite,
- $\forall e \# e'$, $\forall e' \leq_E e''$, $e \# e''$.

We will often omit the indices in \leq_E , $\#_E$ if they are obvious from the context.

In the second axiom, we say that the conflict (e, e'') is **inherited** from (e, e') . If a conflict (e, e') is not inherited (meaning dependencies of e and e' are pairwise **compatible**, *i.e.* non-conflicting), we say that it is a **minimal conflict** and denote it by $e \rightsquigarrow e'$. The *states* of an event structure E , called **configurations**, are the finite sets $x \subseteq E$ that are both consistent (events are pairwise compatible) and **down-closed** (*i.e.* for all $e \in x$, for all $e' \leq e$, then $e' \in x$) – the set of configurations on E is written $\mathcal{C}(E)$, and is partially ordered by inclusion. Configurations with a maximal element are called **prime configurations**, they are those of the form $[e]$ for $e \in E$. We will also use the notation $[e] = [e] \setminus \{e\}$. Between configurations, the **covering relation** $x \prec_C y$ means that y is obtained from x by adding exactly one event: y is an **atomic extension** of x . We might also write $x \xrightarrow{e} y$ to mean that $e \notin x$ and $x \cup \{e\} \in \mathcal{C}(E)$. Finally, when drawing event structures, we will not represent the full partial order \leq but the **immediate causality** generating it, defined as $e \rightarrow e'$ whenever $e < e'$ and for any $e \leq e'' \leq e'$, either $e = e''$ or $e'' = e'$. The diagrams of Figure 1 represent event structures; only displaying immediate causality \rightarrow and minimal conflict \rightsquigarrow .

Arenas. Besides causality and conflict, events in Figure 1 carry *labels* (q, tt, ff, \dots): formally, those will come from the *game*, or the *arena*. Arenas are the semantic representatives of types. They are certain event structures with polarities and Questions/Answers labeling.

► **Definition 2.** An **arena** is a triple $(A, \text{pol}_A, \lambda_A)$ where A is an event structure, $\text{pol}_A : A \rightarrow \{-, +\}$ and $\lambda_A : A \rightarrow \{\mathcal{Q}, \mathcal{A}\}$ are labelings for polarity and Questions/Answers, such that:

- The order \leq_A is *forest-shaped*: for $a_1, a_2 \leq a \in A$, either $a_1 \leq a_2$ or $a_2 \leq a_1$.
- The relation \rightarrow_A is *alternating*: if $a_1 \rightarrow_A a_2$, then $\text{pol}_A(a_1) \neq \text{pol}_A(a_2)$.
- If $\lambda_A(a_2) = \mathcal{A}$, then there is $a_1 \rightarrow a_2$, and $\lambda_A(a_1) = \mathcal{Q}$,
- A is *race-free*: if $a_1 \rightsquigarrow a_2$, then $\text{pol}_A(a_1) = \text{pol}_A(a_2)$.

An arena A is **negative** if all its minimal events have negative polarity.

Conflict aside, our arenas resemble those of [13]: the justification relation traditionally denoted by \vdash_A is simply \rightarrow_A . Basic arenas include the empty arena 1 , and the arena $\llbracket \mathbb{B} \rrbracket$ displayed in Figure 2, often written \mathbb{B} by abuse of notation, for booleans.

We mention here some constructions on arenas. The **dual** A^\perp of A is obtained by taking $\text{pol}_{A^\perp} = -\text{pol}_A$, and leaving the rest unchanged. The **simple parallel composition** $A \parallel B$ is obtained as having events the tagged disjoint union $\{1\} \times A \cup \{2\} \times B$, and all components inherited. The **product** $A \& B$ of negative A and B is obtained as $A \parallel B$, with all events of A in conflict with events of B . As types will be denoted by *negative* arenas, we need a negative arena to interpret \multimap . This is done by setting A to depend on (negative) minimal events of B . If B has at most one minimal event, this is easy:

► **Definition 3.** Let A, B be negative arenas. Assume that B is **well-opened**, *i.e.* $\min(B)$ has at most one event. If $B = 1$, then $A \multimap B = 1$. Otherwise, $\min(B) = \{b_0\}$. We define $A \multimap B$ as $A^\perp \parallel B$, with the additional causal dependency $(2, b_0) \leq (1, a)$ for all $a \in A$.

Defining $A \multimap B$ for well-opened B is sufficient to interpret $\mathbf{aPCF}_{\text{por}}$ types, but would be insufficient in the presence of a tensor type; with *e.g.* $\mathbb{B} \multimap (\mathbb{B} \otimes \mathbb{B})$. The complication here is due to a *disjunctive causality* at the level of types: the left \mathbb{B} is caused by either of the occurrences of \mathbb{B} on the right. As for parallel-or, the inability of event structures to express that can be worked around by introducing two conflicting copies of the left \mathbb{B} , one for each causal justification – this is done in [6], and is reminiscent of the standard arena construction of [13]. This works well for some purposes, but the informed reader may see why this threatens definability for a concurrent language with tensor: indeed a counter-strategy can then behave differently depending on the cause of an occurrence of the left \mathbb{B} .

We avoid the issue here, and only build \multimap for $\mathbf{aPCF}_{\text{por}}$ types. Interestingly the issue vanishes in Section 3: **Disj** supports disjunctive causality in both arenas and strategies.

Prestrategies on arenas. Strategies are certain event structures labeled by arenas. More formally, the labeling function is required to be a *map of event structures*:

► **Definition 4.** A **prestrategy** on arena A is a function on events $\sigma : S \rightarrow A$ *s.t.* σ is a **map of event structures**: it preserves configurations ($\forall x \in \mathcal{C}(S), \sigma x \in \mathcal{C}(A)$) and is locally injective ($\forall s_1, s_2 \in x \in \mathcal{C}(S), \sigma s_1 = \sigma s_2 \implies s_1 = s_2$).

Event structures and their maps form a category \mathcal{E} . Figure 1 displays such prestrategies $\sigma : S \rightarrow A$ – the event structure drawn is S , and events are annotated by their image in A through σ . *Strategies*, introduced in the next section, will be subject to further conditions.

2.3 An interpretation of $\mathbf{aPCF}_{\text{por}}$ as strategies

Following the methodology of denotational semantics, giving an interpretation of $\mathbf{aPCF}_{\text{por}}$ consists in constructing a category out of certain prestrategies, with enough structure to interpret $\mathbf{aPCF}_{\text{por}}$. For lack of space we can only sketch part of the construction, a more detailed account of the construction, originally from [17], can be found in [7].

A **prestrategy from A to B** is a prestrategy $\sigma : S \rightarrow A^\perp \parallel B$. Given also $\tau : T \rightarrow B^\perp \parallel C$, we wish to *compose* them. As usual in game semantics, this involves (1) parallel interaction where the strategies freely communicate, and (2) hiding. We focus on (1) first.

Interaction. The interaction of $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$ is an event structure $T \otimes S$, labeled by $A \parallel B \parallel C$ via $\tau \otimes \sigma : T \otimes S \rightarrow A \parallel B \parallel C$.

The omitted definition [7] of $T \otimes S$ follows the lines of the third author’s event structure semantics for parallel composition in CCS [20]. It is uniquely determined (up to iso) as a *pullback* [7] in \mathcal{E} :

$$\begin{array}{ccccc}
 & & T \otimes S & & \\
 & \Pi_1 \swarrow & & \searrow \Pi_2 & \\
 S \parallel C & & \tau \otimes \sigma & & A \parallel T \\
 & \swarrow \sigma & \downarrow & \swarrow C^A & \searrow \tau \\
 & & \sigma \parallel C \parallel A \parallel B \parallel C^A & &
 \end{array}$$

It is the “smallest” labeled event structure accommodating the constraints of σ and τ .

Hiding. Composition should yield a prestrategy from A to C . Accordingly events of $T \otimes S$ that map to A or C are available to the outside world, and called *visible*. On the other hand, those mapping to B are private synchronization events, dubbed *invisible*; that we wish to *hide*. The proposition below formalizes that event structures are expressive enough for that.

strategies $\llbracket M \rrbracket \in \mathbf{CG}_-(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket)$. Divergence \perp is interpreted as the strategy (also written \perp) with no positive moves, closed by receptivity. Constants \mathbb{t}, \mathbb{f} are the obvious strategies on \mathbb{B} . For $\Gamma \vdash M : \mathbb{B}$, $\Delta \vdash N_1, N_2 : \mathbb{B}$, $\llbracket \text{if } M \text{ then } N_1 \text{ else } N_2 \rrbracket = \mathbf{if} \odot (\llbracket M \rrbracket \otimes \langle \llbracket N_1 \rrbracket, \llbracket N_2 \rrbracket \rangle)$, where \mathbf{if} is given in Figure 3. Finally, parallel-or is the strategy on the right of Figure 1.

The interpretation $\llbracket - \rrbracket$ is adequate: for all $\vdash M : \mathbb{B}$, $M \Downarrow$ iff $\llbracket M \rrbracket \neq \perp$. However, we will be better equipped to prove that in Section 4. For now, the desired induction invariant for soundness $M \Downarrow v \Rightarrow \llbracket M \rrbracket = \llbracket v \rrbracket$ fails, as the model keeps track of too much intensional information: $\llbracket \text{por } \mathbb{t} \ \mathbb{t} \rrbracket$ has two conflicting \mathbb{t}^+ events, and is therefore not isomorphic to $\llbracket \mathbb{t} \rrbracket$.

2.4 Disjunctive causality and observational determinism

We have seen that \mathbf{CG}_- is an adequate model of $\mathbf{aPCF}_{\text{por}}$. However, it is not *intensionally fully abstract*: strategies distinguish more than terms of $\mathbf{aPCF}_{\text{por}}$. In fact, as in [6] one can interpret *e.g.* linearly used boolean references in \mathbf{CG}_- ; those can obviously distinguish more than the input-output behaviour of terms. We can remove these by requiring conditions of visibility, well-bracketing and innocence as in [8]; and we will do so in Section 4.2.


One condition of [8] is however inadequate: *determinism*. Indeed, the strategy of Figure 1 is non-deterministic, and no deterministic strategy can implement parallel-or (indeed, all deterministic strategies on first-order types implement sequential functions [8]). Formalizing a notion of *observational determinism*, accepting the strategy for parallel-or but rejecting genuinely non-deterministic strategies, proved very challenging.

Disjunctive causality. As we pointed out in the introduction, the non-determinism of parallel-or comes from the incapacity of our version of event structures, sometimes referred to as *prime* event structures for disambiguation, to express *disjunctive causality*. So the reader may wonder: why do we bother with prime event structures at all? After all, there are alternatives. *General event structures* [19] allow events to be enabled in several distinct ways. We give their equivalent presentation [21] in terms of *configuration families*:



► **Definition 9.** A **configuration family (cf)** is a pair $(|\mathcal{A}|, \mathcal{A})$ (often just written \mathcal{A}) where $|\mathcal{A}|$ is a set of events, and \mathcal{A} is a set of **configurations**, which are finite subsets of $|\mathcal{A}|$, such that $\emptyset \in \mathcal{A}$, and satisfying the two further axioms:

- *Completeness.* For $x, y \in \mathcal{A}$, if $x \uparrow y$ (they are **compatible**, *i.e.* there exists $z \in \mathcal{A}$ such that $x \subseteq z$ and $y \subseteq z$), then $x \cup y \in \mathcal{A}$.
- *Coincidence-freeness.* If $x \in \mathcal{A}$, for all distinct $e_1, e_2 \in x$ there exists $y \subseteq x$ such that $y \in \mathcal{A}$ and $e_1 \in y \Leftrightarrow e_2 \notin y$.

For every (prime) event structure A , $\mathcal{C}(A)$ is a configuration family. However configuration families are more general. In a configuration family, the same event can occur for different reasons. For instance, on the set of events $\{\ominus_1, \ominus_2, \oplus\}$, Figure 4 represents a configuration family where the event \oplus can occur either because of \ominus_1 or \ominus_2 , represented symbolically as in the diagram on the right – importantly, an occurrence of \oplus carries no data on its effective cause. With the same symbolic representation, the causally disjunctive nature of parallel-or is made explicit in Figure 5. It seems clear that this is a more accurate description of parallel-or. In particular, it is *deterministic*, in a sense to be made formal in the next section.



Are configuration families sufficiently expressive as a basis to construct a category of strategies, as we did with prime event structures in the previous section? Recall that for that we used two properties of event structures: that the corresponding category \mathcal{E} has *pullbacks*




■ **Figure 4** Disj. caus. in configuration families

■ **Figure 5** por as a disj. strategy

(for interactions), and then that event structures support *hiding*. We will show very soon that the first criterion holds. However, we run into an issue for the second:

► **Proposition 10.** Configuration families do *not* support hiding. There is a cf \mathcal{A} and $V \subseteq |A|$ such that the set $\{x \cap V \mid x \in \mathcal{A}\}$ is not a configuration family.

Proof. Set $|A| = \{a, b, c, d\}$, and configurations those specified by: (d is caused either by a , or by b and c together). With $V = \{b, c, d\}$, the obtained hidden set fails completeness: it contains $\{b\}$, $\{d\}$, $\{b, c, d\}$ but not $\{b, d\}$. ◀



It is part for this reason that prime event structures are used when building categories of games and strategies – as indeed, hiding is crucial. Removing either coincidence-freeness or completeness loses the correspondence with general event structures [21], and leads to various pathologies further down the road – for instance we lose the key lemma:

► **Lemma 11.** Let \mathcal{A} a set of configurations on $|A|$ satisfying completeness. Then, it is coincidence-free iff for any $x, y \in \mathcal{A}$ s.t. $x \subseteq y$, there is a covering chain in \mathcal{A} : $x \subsetneq \dots \subsetneq y$.

3 Disjunctive deterministic games

In this section we introduce the main contribution of the paper, a category **Disj** of disjunctive deterministic strategies supporting the interpretation of parallel-or. This relies on a notion of *deterministic* configuration families, and the observation that those *do* support hiding.

3.1 Deterministic configuration families

First we adjoin cfs with *polarities*, and define determinism (for Player).

► **Definition 12.** A cf with polarities (cfp) (*resp.* partial polarities (cfpp)) is a cf \mathcal{A} with $\text{pol}_{\mathcal{A}} : |A| \rightarrow \{-, +\}$ (*resp.* $\text{pol}_{\mathcal{A}} : |A| \rightarrow \{-, 0, +\}$).

A cfpp \mathcal{A} is **deterministic** (dcfp/dcfpp) iff for all $x \subseteq_{\mathcal{A}}^p y$ and $x \subseteq z$, we have $y \cup z \in \mathcal{A}$ – we write $x \subseteq_{\mathcal{A}}^p y$ (*resp.* $x \subseteq_{\mathcal{A}}^- y$, $x \subseteq_{\mathcal{A}}^+ y$) to mean that $x \subseteq y$ and $\text{pol}(y \setminus x) \subseteq \{0, +\}$ (*resp.* $\text{pol}_{\mathcal{A}}(y \setminus x) \subseteq \{-\}$, $\text{pol}_{\mathcal{A}}(y \setminus x) \subseteq \{+\}$).

This means that only Opponent makes irreversible choices regarding the evolution of the play. We show that dcfpps support hiding of neutral events. In the sequel, given a dcfpp \mathcal{A} , V always denotes the set of non-neutral events in $|A|$. We write \mathcal{A}_{\downarrow} for the candidate cf on events V with configurations those of the form $x_{\downarrow} = x \cap V$ for $x \in \mathcal{A}$. The following key lemma is proved by successive applications of determinism and completeness.

► **Lemma 13.** For $x \in \mathcal{A}_{\downarrow}$, $y \in \mathcal{A}$ such that $x \subseteq y_{\downarrow}$, for any $x' \in \mathcal{A}$ a witness for x (i.e. $x'_{\downarrow} = x$), there is $y \subseteq y' \in \mathcal{A}$ such that $x' \subseteq y'$.

From Lemma 13 follows that for $x, y \in \mathcal{A}$, if $x_{\downarrow} \uparrow y_{\downarrow}$ in \mathcal{A}_{\downarrow} , then $x \uparrow y$ in \mathcal{A} . Using that observation together with Lemma 13 and determinism, it is straightforward to prove:

► **Proposition 14.** For \mathcal{A} a dcfpp, \mathcal{A}_{\downarrow} is a dcfp.

3.2 Deterministic disjunctive strategies and composition

A **pregame** is simply a cfp – for \mathcal{A}, \mathcal{B} cfps, \mathcal{A}^{\perp} has the same events and configurations as \mathcal{A} but inverted polarity. We will also write \mathcal{B}^0 for the cfpp with the same configurations as \mathcal{B} but polarity set as globally 0. The cf $\mathcal{A} \parallel \mathcal{B}$ has events the tagged disjoint union, and configurations $x_A \parallel x_B$ for $x_A \in \mathcal{A}$, $x_B \in \mathcal{B}$; $\mathcal{A} \& \mathcal{B}$ has the same events as $\mathcal{A} \parallel \mathcal{B}$, but only those configurations with one side empty.

Unlike in Definition 4, *disjunctive prestrategies* are simply substructures of the (pre)games.

► **Definition 15.** A **prestrategy** on \mathcal{A} is a dcfp σ on $|\mathcal{A}|$, such that $\sigma \subseteq \mathcal{A}$ – written $\sigma : \mathcal{A}$.

► **Example 16.** On events (with polarities) $\{q^-, tt^+, ff^+\}$, $\mathcal{C}(\mathbb{B})$ is a cfp – that by abuse of notation we still write \mathbb{B} . Then, Figure 5 denotes a prestrategy on $\mathbb{B}^{\perp} \parallel \mathbb{B}^{\perp} \parallel \mathbb{B}$.

Though disjunctive prestrategies are not primarily defined as maps, it will be helpful in composing them that they *can* be. Given two cfs \mathcal{A} and \mathcal{B} , a **map from \mathcal{A} to \mathcal{B}** is a function on events $f : |\mathcal{A}| \rightarrow |\mathcal{B}|$ which preserves configurations and is locally injective. Configuration families and their maps forms a category **Fam**. Note that if A is an event structure, $\mathcal{C}(A)$ is a configuration family on $|A|$. The definition of maps of cfs is compatible with that of maps of event structures, making $\mathcal{C}(-) : \mathcal{E} \rightarrow \mathbf{Fam}$ a full and faithful functor. Finally, a prestrategy $\sigma : \mathcal{A}$ on \mathcal{A} can be regarded as an identity-on-events **Fam**-morphism $\sigma \rightarrow \mathcal{A}$.

Like \mathcal{E} , **Fam** has pullbacks. The **interaction** of σ on $\mathcal{A}^{\perp} \parallel \mathcal{B}$ and τ on $\mathcal{B}^{\perp} \parallel \mathcal{C}$ is the pullback of identity-on-events maps $\sigma \parallel \mathcal{C} \rightarrow \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$ and $\mathcal{A} \parallel \tau \rightarrow \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$.

► **Proposition 17.** The pullback above is (up to iso) the cf $\tau \otimes \sigma$ with events $|\mathcal{A}| \parallel |\mathcal{B}| \parallel |\mathcal{C}|$ and configurations those $x \in (\sigma \parallel \mathcal{C}) \cap (\mathcal{A} \parallel \tau)$ that are **secured**: for distinct $p_1, p_2 \in x$ there is $x \supseteq y \in (\sigma \parallel \mathcal{C}) \cap (\mathcal{A} \parallel \tau)$ s.t. $p_1 \in y \Leftrightarrow p_2 \notin y$.

Equivalently, $x \in (\sigma \parallel \mathcal{C}) \cap (\mathcal{A} \parallel \tau)$ is secured iff it has a *covering chain*, i.e. a sequence: $\emptyset = x_0 - \subset x_1 - \subset \dots - \subset x_n = x$ where for all $0 \leq i \leq n$, $x_i \in (\sigma \parallel \mathcal{C}) \cap (\mathcal{A} \parallel \tau)$.

We regard $\tau \otimes \sigma$ as a cfpp by setting as polarities those of $\mathcal{A}^{\perp} \parallel \mathcal{B}^0 \parallel \mathcal{C}$. To use Proposition 14 and finish defining composition, $\tau \otimes \sigma$ needs to be deterministic; a sufficient condition for that is that σ and τ are *receptive*. A prestrategy σ on \mathcal{A} is **receptive** iff for all $x \in \sigma$, if $x \cup \{a^-\} \in \mathcal{A}$ then $x \cup \{a^-\} \in \sigma$.

► **Proposition 18.** If σ and τ are receptive prestrategies, then $\tau \otimes \sigma$ is a *deterministic* cfpp. Then, $\tau \odot \sigma = (\tau \otimes \sigma)_{\downarrow}$ is a receptive prestrategy on $\mathcal{A}^{\perp} \parallel \mathcal{C}$.

Hence, composition is well-defined on receptive prestrategies – it is also associative. To get a category, we now prove the copycat strategy to be an identity.

3.3 Copycat and the compact closed category Disj



We cannot replicate in pregames the definition of copycat sketched in Section 2.3. However, as observed in [18, 7], if A is an arena, the *configurations* of \mathbb{C}_A are those configurations of $A^{\perp} \parallel A$, necessarily of the form $x_l \parallel x_r$, such that every positive event of x_r (*w.r.t.* A) is already in x_l , and every positive event in x_l (*w.r.t.* A^{\perp}) is already in x_r . In other words, $x_l \supseteq_A^- x_l \cap x_r \subseteq_A^+ x_r$, written $x_r \sqsubseteq_A x_l$ in [18, 7] and referred to as the “Scott order”.

Accordingly, on a pregame \mathcal{A} , given $x, y \in \mathcal{A}$ we write $x \sqsubseteq_{\mathcal{A}} y$ iff $x \cap y \in \mathcal{A}$ and the relation above holds. The candidate prestrategy **copycat** $\alpha_{\mathcal{A}}$ comprises all $x \parallel y \in \mathcal{A}^{\perp} \parallel \mathcal{A}$ s.t. $y \sqsubseteq_{\mathcal{A}} x$. Indeed $\alpha_{\mathcal{A}}$ is a configuration family and is receptive; but prestrategies must be *deterministic*. It turns out that as in [22], $\alpha_{\mathcal{A}}$ is only deterministic when \mathcal{A} is *race-free*:

► **Proposition 19.** Let \mathcal{A} be a pregame. Then, $\alpha_{\mathcal{A}}$ is a prestrategy iff \mathcal{A} is **race-free**: for all $x, y, z \in \mathcal{A}$ such that $x \sqsubseteq_{\mathcal{A}}^+ y$ and $x \sqsubseteq_{\mathcal{A}}^- z$, we have $y \uparrow_{\mathcal{A}} z$.

We aim to reproduce Theorem 6 in this new setting, and characterise the prestrategies left invariant under composition with copycat. However, there is a new subtlety here: for arbitrary \mathcal{A} , $\alpha_{\mathcal{A}}$ might not even be idempotent!

► **Example 20.** Consider the cfp on events $A = \{\ominus_1, \oplus_2, \oplus_3\}$ given on the right. This is a race-free pregame, so by Proposition 19, $\alpha_{\mathcal{A}} \odot \alpha_{\mathcal{A}}$ is a prestrategy on $\mathcal{A}^{\perp} \parallel \mathcal{A}$. However, it is *distinct from* $\alpha_{\mathcal{A}}$. Indeed, the following configuration of $\mathcal{A}^{\perp} \parallel \mathcal{A}^0 \parallel \mathcal{A}$ belongs to $\alpha_{\mathcal{A}} \otimes \alpha_{\mathcal{A}}$:



After hiding, there is a change in the causal history of \oplus_3 , not authorized by $\alpha_{\mathcal{A}}$ but authorised by \mathcal{A} : it is caused by \ominus_2 on the left, but \ominus_1 on the right. This issue comes from the fact that in \mathcal{A} , the same event can be caused either by a positive or a negative move. Such behaviour will be banned in games:

► **Proposition 21.** If \mathcal{A} is a race-free pregame, the following are equivalent: (1) $\alpha_{\mathcal{A}}$ is idempotent, (2) $\sqsubseteq_{\mathcal{A}}$ is a partial order, (3) \mathcal{A} is **co-race-free**: for all $x, y, z \in \mathcal{A}$ with $x \supseteq^- y, x \supseteq^+ z$, we have $y \cap z \in \mathcal{A}$.

A **game** will be a race-free, co-race-free pregame. Copycat on any game is an idempotent prestrategy – *strategies* are those prestrategies that compose well with copycat. We prove:

► **Theorem 22.** Let σ be a receptive prestrategy on $\mathcal{A}^{\perp} \parallel \mathcal{B}$. Then, $\alpha_{\mathcal{B}} \odot \sigma \odot \alpha_{\mathcal{A}} = \sigma$ iff σ is **courteous**: for any $x \dashv x \cup \{a_1^+\} \dashv x \cup \{a_1^+, a_2^+\}$ in σ , if $x \cup \{a_2^+\} \in \mathcal{A}^{\perp} \parallel \mathcal{B}$, then $x \cup \{a_2^+\} \in \sigma$ as well. In this case, σ is a **strategy**, written: $\sigma : \mathcal{A}^{\perp} \parallel \mathcal{B}$.

It follows that there is a category **Disj** with games as objects, and strategies $\sigma : \mathcal{A}^{\perp} \parallel \mathcal{B}$ as morphisms from \mathcal{A} to \mathcal{B} . It is fairly easy to show that just as **CG**, this category is compact closed. But unlike **CG**, **Disj** supports a deterministic interpretation of parallel-or: indeed the set of configurations of $\mathbb{B}^{\perp} \parallel \mathbb{B}^{\perp} \parallel \mathbb{B}$ denoted by the diagram of Figure 5 is a strategy.

3.4 An SMCC and deterministic interpretation of $\mathbf{aPCF}_{\text{por}}$

As for **CG**, **Disj** lacks structure to interpret $\mathbf{aPCF}_{\text{por}}$: the family of bottom strategies $e_{\mathcal{A}} : \mathcal{A}^{\perp} \parallel 1$ (where again 1 is the empty game) fails naturality. As before, we hence restrict **Disj** to a subcategory of *negative* games and strategies.

► **Definition 23.** A cfp \mathcal{A} is **negative** if any non-empty $x \in \mathcal{A}$ includes a negative event.

Copycat on negative \mathcal{A} is negative and negative strategies are stable under composition; so there is a subcategory **Disj₋** of **Disj** with negative games as objects and negative strategies as morphisms, inheriting a symmetric monoidal structure from **Disj**. Moreover 1 is terminal, and **Disj₋** has products given by $\mathcal{A} \& \mathcal{B}$.

To prove the monoidal closure, as in **CG**, we have to deal with the fact that for \mathcal{A} and \mathcal{B} negative, $\mathcal{A}^{\perp} \parallel \mathcal{B}$ is not necessarily negative, so we define:

► **Definition 24.** Let \mathcal{A} and \mathcal{B} be two negative games. The game $\mathcal{A} \multimap \mathcal{B}$ has the same events (and polarity) as $\mathcal{A}^\perp \parallel \mathcal{B}$, but non-empty configurations those $x_A \parallel x_B \in \mathcal{A}^\perp \parallel \mathcal{B}$ such that x_B is non-empty (and hence includes a negative event).

Recall from Definition 3 that the arrow arena $A \multimap B$ was only defined for B *well-opened* (i.e. with at most one minimal event). This was to avoid constructing arenas for types like $\mathbb{B} \multimap (\mathbb{B} \otimes \mathbb{B})$ (invalid for $\mathbf{aPCF}_{\text{por}}$, but valid in an extension with a tensor type), where the left hand side \mathbb{B} can be opened because of either of the right hand side occurrences of \mathbb{B} . In **Disj**, when constructing $\mathcal{A} \multimap \mathcal{B}$ there is no well-openness condition on \mathcal{B} : exploiting disjunctive causality we can express that \mathcal{A} is opened only after *some* event of \mathcal{B} – it does not matter which one. The negative strategies in $\mathcal{A}^\perp \parallel \mathcal{B}^\perp \parallel \mathcal{C}$ and $\mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C}$ are the same, from which (using the compact closed structure of **Disj**) it follows that **Disj**₋ is monoidal closed, instead of only having an exponential ideal. Like **CG**₋, it supports an adequate interpretation of $\mathbf{aPCF}_{\text{por}}$. Unlike **CG**₋, its strategies are *deterministic*.

4 Glueing **CG**₋ and **Disj**₋, and full abstraction

By moving from **CG**₋ to **Disj**₋, we gain determinism. However, **Disj**₋ is *not* intensionally fully abstract – there are some strategies that distinguish syntactically undistinguishable terms of $\mathbf{aPCF}_{\text{por}}$. For instance, one has a strategy on $(\mathbb{B} \multimap \mathbb{B} \multimap \mathbb{B}) \multimap \mathbb{B}$ that calls its argument, feeds it tt as first argument, and tt as second argument only if the first argument has been evaluated; it then copies the final result to toplevel. Doing so, it distinguishes the observationally equivalent $\lambda xy. \text{if } x \text{ then (if } y \text{ then } \text{tt} \text{ else } \perp) \text{ else } \perp$ and $\lambda xy. \text{if } y \text{ then (if } x \text{ then } \text{tt} \text{ else } \perp) \text{ else } \perp$.

Getting rid of such strategies is the responsibility of the concept of *innocence* in Hyland-Ong games [13]. But the *P-views* of innocent strategies carry precisely the *causal* information that we have lost when moving from **CG**₋ to **Disj**₋! This causal information was crucial in [8] to give concurrent versions of well-bracketing and innocence, so as to capture the behaviour of parallel **PCF** strategies. So, the deterministic account of $\mathbf{aPCF}_{\text{por}}$ is not enough; the causal (where innocence and well-bracketing are defined) and deterministic models should be related, to establish in what sense the causal model is already “observably” deterministic.

4.1 A glued games model

For any arena A (Definition 2), $\mathcal{C}(A)$ is a game – it is race-free and co-race-free. Moreover, arena constructions and game constructions match: for any A, B , $\mathcal{C}(A^\perp) = \mathcal{C}(A)^\perp$, $\mathcal{C}(A \parallel B) = \mathcal{C}(A) \parallel \mathcal{C}(B)$; for negative A, B we have $\mathcal{C}(A \& B) = \mathcal{C}(A) \& \mathcal{C}(B)$, and if B is well-opened, $\mathcal{C}(A \multimap B) = \mathcal{C}(A) \multimap \mathcal{C}(B)$. Therefore, we will just take the objects of our glued games model to be arenas. The strategies, though, will be strategies in *both categories*.

► **Definition 25.** Let A be an arena. A strategy $\sigma : S \rightarrow A$ is **observationally deterministic (odet)** if (1) the **display** of σ , $\mathcal{O}(\sigma) = \{\sigma x \mid x \in \mathcal{C}(S)\}$ is a disjunctive deterministic strategy on $\mathcal{C}(A)$ (in **Disj**). It follows then that σ is also a **Fam**-morphism $\sigma : \mathcal{C}(S) \rightarrow \mathcal{O}(\sigma)$. We also ask that it has (2) the **configuration extension property**: for any $x \in \mathcal{C}(S)$, if $\sigma x \subseteq y \in \mathcal{O}(\sigma)$, then there is $x' \in \mathcal{C}(S)$ such that $\sigma x' = y$.

The configuration extension property ensures that two causal realizations of the same displayed configuration have bisimilar futures. The display operation yields a coarser equivalence on strategies: odet strategies $\sigma, \tau : A$, are **display-equivalent**, written $\sigma \approx \tau$, if $\mathcal{O}(\sigma) = \mathcal{O}(\tau)$. This coarser equivalence is key to ensure soundness of our interpretation, as the interpretation in **CG**₋ of $\text{por } \text{tt } \text{tt}$ and tt are *not* isomorphic, but only display-equivalent.

► **Theorem 26.** *There is a compact closed category \mathbf{Odet} with arenas as objects and odet strategies up to \approx as morphisms; with a symmetric monoidal subcategory \mathbf{Odet}_- of negative arenas and negative strategies with products and 1 terminal. It has well-opened arenas as an exponential ideal. Finally, \mathbf{Odet}_- supports an adequate interpretation of $\mathbf{aPCF}_{\text{por}}$.*

Proof. For odet σ and τ , $\mathcal{O}(\tau \odot \sigma) = \mathcal{O}(\tau) \odot \mathcal{O}(\sigma)$: \subseteq is direct, \supseteq exploits the extension property for σ and τ . So, $\mathcal{O}(\tau \odot \sigma)$ is a deterministic disjunctive strategy, and it has the extension property – hence $\tau \odot \sigma$ is odet. In general, \mathcal{O} links the structure of \mathbf{CG} to that of \mathbf{Disj} , forming \mathbf{Odet} . We sketch adequacy.

Soundness. For all $M \Downarrow v$, $\llbracket M \rrbracket = \llbracket v \rrbracket$ (by induction on the evaluation derivations).

Adequacy. We prove by induction on the size of M that if $M \Uparrow$, then $\llbracket M \rrbracket$ is bottom. The size is an adequate measure because as $\mathbf{aPCF}_{\text{por}}$ is affine, substitution is non-copying and each induction step is on a strictly smaller term. ◀

As a result, interpretations in \mathbf{CG}_- and \mathbf{Disj} are also adequate. We have accommodated the causal representation of programs permitted by \mathbf{CG} and the determinism of \mathbf{Disj} . Now, it remains to import the causal conditions on strategies of [8], and prove full abstraction.


4.2 Conditions

We now recall *innocence* and *well-bracketing*, introduced in [8]. Formulated for *deterministic strategies* (in the sense of \mathbf{CG}_-), those would not suffice to prove full abstraction for a genuinely non-deterministic language. For that, further conditions are needed to ensure the locality of conflicts – those will appear in the first author’s forthcoming PhD thesis. However, we will show in Lemma 31 that for $\mathbf{aPCF}_{\text{por}}$, distinguishable strategies can be distinguished via *deterministic* (in the sense of \mathbf{CG}_-) contexts, so these simpler conditions will suffice.

All our conditions rely crucially on the notion of *grounded causal chain*.

► **Definition 27.** A **grounded causal chain (gcc)** in S is a sequence $\rho = \rho_0 \rightarrow \dots \rightarrow \rho_n$ where ρ_0 is minimal. The set of gccs of S is written $\text{gcc}(S)$.

By courtesy, gccs of strategies on arenas are always alternating. They give a notion of *thread* of a concurrent strategy. Two gccs $\rho, \rho' \in \text{gcc}(S)$ are **forking** when $\rho \cup \rho'$ is consistent in S , and there is $k \in \mathbb{N}$ s.t. $\rho_i = \rho'_i$ for $i < k$ and $\{\rho_i\}_{i \geq k}$ and $\{\rho'_j\}_{j \geq k}$ are non-empty and disjoint. They are **negatively forking** when $\text{pol}(\rho_{k+1}) = \text{pol}(\rho'_{k+1}) = -$, **positively forking** otherwise. Two forking gccs ρ, ρ' are **joined** at $s \in S$ when $\rho_\omega \rightarrow s$ and $\rho'_\omega \rightarrow s$. (ρ_ω refers to the last event of ρ) as shown in this picture:



Innocence. *Innocence* enforces independence between threads forked by Opponent.

► **Definition 28.** Let $\sigma : S \rightarrow A$ be a strategy on a arena A . It is **innocent** when it is:

Visible: The image $\sigma \rho$ of a gcc (regarding $\rho \in \text{gcc}(S)$ as a set) is a configuration of A .

Preinnocent: Two negatively forking gccs are never joined.

Preinnocence is a *locality* condition, forcing Player to deal independently with threads forked by Opponent – in the sequential case, it coincides with Hyland-Ong innocence. In the concurrent case though, it is not by itself stable under composition; that is where visibility comes in. Together they are preserved under composition and under all the operations on strategies involved in the interpretation of $\mathbf{aPCF}_{\text{por}}$ in \mathbf{CG}_- (and hence \mathbf{Odet}_-).

Well-bracketing. Traditionally, *well-bracketing* in game semantics rules out strategies that behave like a control operator, manipulating the call stack. It exploits the question/answer labelling, reminiscent of the function calls/returns. In arenas from $\mathbf{aPCF}_{\text{por}}$ types, answers to the same question are always conflicting, so every consistent set has at most one answer to any question. A consistent set X is **complete** when it has exactly one answer to any question. A question is **pending** in a set X if it has no answer in X and maximally so in X .

► **Definition 29.** A strategy $\sigma : S \rightarrow A$ is **well-bracketed** when (1) if $a \in S$ answers $q \in S$, then q is pending in $[a]$ and (2) for $\rho, \rho' \in \text{gcc}(S)$ forking at $\rho_k = \rho'_k$, and joined, the segments $\rho_{>k}$ and $\rho'_{>k'}$ must be complete.

The affinity condition of [8] comes for free here, thanks to the conflict in \mathbb{B} . Well-bracketing is proved stable under composition and other operations in [8]. A $\mathbf{aPCF}_{\text{por}}$ -strategy is a negative, innocent, well-bracketed and odet strategy.

► **Theorem 30.** *There is a symmetric monoidal category $\mathbf{PorStrat}$ of negative arenas and $\mathbf{aPCF}_{\text{por}}$ -strategies, with products, a terminal object and an exponential ideal comprising \mathbb{B} . Moreover, the interpretation of $\mathbf{aPCF}_{\text{por}}$ in \mathbf{Odet}_- factors through $\mathbf{PorStrat} \subseteq \mathbf{Odet}_-$.*

From now on, all strategies will be assumed to be $\mathbf{aPCF}_{\text{por}}$ -strategies.

4.3 Intensional full abstraction

Two strategies $\sigma, \tau : A$ are **observationally equivalent** ($\sigma \simeq_{\text{obs}} \tau$) when for all strategies $\alpha : A^\perp \parallel \mathbb{B}$, $\alpha \odot \sigma \approx \alpha \odot \tau$. In fact, there is no need to quantify over *all* strategies. A **path-strategy** is a strategy $\sigma : S \rightarrow A$ such that there exists a configuration $x \in \mathcal{C}(S)$ that contains all the positive moves of S . Any distinguishing strategy yields by restriction a distinguishing path-strategy – this would fail without affinity, as the restriction would fail uniformity [8], and enforcing uniformity would make it non-deterministic.

► **Lemma 31.** *Two distinguishable strategies can be distinguished by a path-strategy.*

► **Lemma 32.** *Every path-strategy can be defined by a term of \mathbf{aPCF} up to \simeq_{obs} .*

Proof. As in [8] since path-strategies are deterministic – the only difference is the necessity to distribute the context among the subterms to ensure affine typing. ◀

A corollary is that in an affine setting, **por** adds no distinguishing power: two $\mathbf{aPCF}_{\text{por}}$ terms are observationally equivalent if and only if they cannot be distinguished by a context from affine \mathbf{PCF} without **por**. Intensional full abstraction follows by standard techniques:

► **Theorem 33.** *The interpretation of $\mathbf{aPCF}_{\text{por}}$ into $\mathbf{PorStrat}$ is intensionally fully abstract: it preserves and reflects observational equivalence (see Section 2.1).*

5 Conclusion

This leaves many avenues for further investigation. On the semantic front, we plan among other applications to exploit \mathbf{Odet} to model non-interfering concurrent languages. On the foundational front we need to understand better how the present approach relates to the treatment of disjunctive causality in edcs [9].

Acknowledgments. We are grateful to Tamás Kispéter for interesting discussions about strategies as general event structures.

References

- 1 Samson Abramsky, Kohei Honda, and Guy McCusker. A fully abstract game semantics for general references. In *LICS, Indianapolis, Indiana, USA, June 21-24, 1998*, pages 334–344. IEEE Computer Society, 1998.
- 2 Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.
- 3 Samson Abramsky and Guy McCusker. Linearity, sharing and state: a fully abstract game semantics for idealized algol with active expressions. *ENTCS*, 3:2–14, 1996.
- 4 Samson Abramsky and Paul-André Melliès. Concurrent games and full completeness. In *LICS, Trento, Italy, July 2-5, 1999*, pages 431–442. IEEE Computer Society, 1999.
- 5 Robert Cartwright, Pierre-Louis Curien, and Matthias Felleisen. Fully abstract semantics for observably sequential languages. *Inf. Comput.*, 111(2):297–401, 1994.
- 6 Simon Castellan and Pierre Clairambault. Causality vs. interleavings in concurrent game semantics. In *CONCUR 2016, August 23-26, 2016, Québec City, Canada*, pages 32:1–32:14.
- 7 Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. 2016. Submitted, <https://arxiv.org/abs/1604.04390>.
- 8 Simon Castellan, Pierre Clairambault, and Glynn Winskel. The parallel intensionally fully abstract games model of PCF. In *LICS, Kyoto, Japan, July 6-10, 2015*, pages 232–243. IEEE Computer Society, 2015.
- 9 Marc de Visme and Glynn Winskel. Strategies with parallel causes. To appear in CSL 2017.
- 10 Dan R. Ghica. Applications of game semantics: From program analysis to hardware synthesis. In *LICS, 11-14 August 2009, Los Angeles, CA, USA*, pages 17–26. IEEE Computer Society, 2009.
- 11 Dan R. Ghica and Andrzej S. Murawski. Angelic semantics of fine-grained concurrency. *Ann. Pure Appl. Logic*, 151(2-3):89–114, 2008.
- 12 Tom Hirschowitz. Full abstraction for fair testing in CCS (expanded version). *Logical Methods in Computer Science*, 10(4), 2014.
- 13 J. M. E. Hyland and C.-H. Luke Ong. On full abstraction for PCF: I, II, and III. *Inf. Comput.*, 163(2):285–408, 2000.
- 14 James Laird. Full abstraction for functional languages with control. In *Proceedings, 12th Annual IEEE Symposium on Logic in Computer Science, Warsaw, Poland, June 29 - July 2, 1997*, pages 58–67. IEEE Computer Society, 1997.
- 15 Paul-André Melliès and Samuel Mimram. Asynchronous games: Innocence without alternation. In *CONCUR, Lisbon, Portugal, September 3-8, 2007*, volume 4703 of *Lecture Notes in Computer Science*, pages 395–411. Springer, 2007.
- 16 Gordon D. Plotkin. LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–255, 1977.
- 17 Silvain Rideau and Glynn Winskel. Concurrent strategies. In *LICS, June 21-24, 2011, Toronto, Ontario, Canada*, pages 409–418, 2011.
- 18 Glynn Winskel. Strategies as profunctors. In *FOSSACS, Held as Part of ETAPS 2013, Rome, Italy, March 16-24, 2013*, pages 418–433.
- 19 Glynn Winskel. *Events in computation*. PhD Thesis, Edinburgh University, 1980.
- 20 Glynn Winskel. Event structure semantics for CCS and related languages. In *ICALP, Aarhus, Denmark, July 12-16, 1982, Proceedings*, pages 561–576, 1982.
- 21 Glynn Winskel. Event structures. In *Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part II, Proceedings of an Advanced Course, Bad Honnef, 8.-19. September 1986*, pages 325–392, 1986.
- 22 Glynn Winskel. Deterministic concurrent strategies. *Formal Asp. Comput.*, 24(4-6):647–660, 2012.