



HAL
open science

Affine strategies in arena games

Russ Harmer

► **To cite this version:**

Russ Harmer. Affine strategies in arena games. 1st International Workshop on Games for Logic and Programming Languages, 2005, Edinburgh, United Kingdom. hal-00013015

HAL Id: hal-00013015

<https://hal.science/hal-00013015>

Submitted on 2 Nov 2005

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Affine strategies in arena games

Russ Harmer

CNRS & Université de Paris 7

Abstract

We show how to extract an SMCC of arenas and affine strategies from the CCC of arenas and innocent strategies, a process that essentially reverses the more usual construction of a CCC from an SMCC and the ! of linear logic.

1 Introduction

Arena-based game semantics [1, 4, 5, 7] has proved highly effective for constructing fully abstract models for a wide variety of sequential programming languages. In contrast to many other game models (AJM, polarized, tree-based, etc) its precise relation with linear logic remains rather unclear: one typically considers a category of innocent (or single-threaded) strategies which turns out to be a CCC with no further ado, in particular having no need for the ! and & of linear logic to get a genuine categorical product [6].

In this paper, we show how to isolate a “lluf” subcategory of strategies which behave in an “affine” manner. As well as clarifying the relation of arena games to linear logic, this could prove useful for developing models of SCI (which typically begin with the affine typing discipline of basic SCI) and also provides a new axis (alongside the established notions of bracketing and rigidity) for analysing the behaviour of innocent strategies. Most importantly, this work marks the first step of an analysis of *innocent interaction*, focusing on a special case which somewhat tames the combinatorial complexity of the general case.

2 Innocent game semantics

Let Σ be a countable set. A **pointing string** over Σ is a string $s \in \Sigma^*$ with pointers between the occurrences of s such that pointers always point to *earlier* occurrences and we have *at most one* pointer from each occurrence. We write s_i for the i th symbol of s (assuming it has one), s_ω for the last occurrence of s (again, if it has one), $|s|$ for the length of s and ε for the empty pointing string (of length 0).

We write $s \sqsubseteq t$ obvious extension of the prefix partial ordering to pointing strings and $s \wedge t$ for the *longest common prefix* of s and t . Finally, if s is a pointing string over Σ and $a \in \Sigma$ then we write $s \cdot a$ for the pointing string obtained by adding a to the end of s , pointing to the *last* occurrence of s .

If $\Sigma' \subseteq \Sigma$, we write $s \upharpoonright \Sigma'$ for the restriction of s to Σ' , *i.e.* that pointing string obtained by removing all symbols of s not from Σ' , keeping pointers between occurrences from Σ' and adding a pointer from s_i to s_j iff following pointers from s_i leads directly into the “forbidden zone” $\Sigma - \Sigma'$ and remains there until reemerging at s_j .

2.1 Arenas, plays and strategies

Game semantics posits two protagonists, **Opponent** and **Player**, that can take turns to play certain “moves” according to the rules of a “game”. We formalize the ideas of “the rules of a game” and “play” as *arenas* and the *legal plays* of a given arena.

An **arena** A is a tuple $\langle M_A, \lambda_A, I_A, \vdash_A \rangle$ where

- M_A is a countable set of **tokens**.
- $\lambda_A : M_A \rightarrow \{\mathbf{O}, \mathbf{P}\} \times \{\mathbf{Q}, \mathbf{A}\}$ labels each $m \in M_A$ as belonging to either Opponent or Player and as being a *question* or an *answer*. We write $\lambda_A^{\mathbf{OP}}$ (resp. $\lambda_A^{\mathbf{QA}}$) for the composite with first (resp. second) projection and $\lambda_A^{\mathbf{PO}}$ for the “inverted” labelling (exchange of \mathbf{O} and \mathbf{P}).
- I_A is a subset of $\lambda_A^{-1}(\mathbf{OQ})$, the **initial tokens** of A .
- \vdash_A is a binary relation on M_A , known as **enabling**, which must satisfy
 - (e1) if $m \vdash_A n$ then $\lambda_A^{\mathbf{OP}}(m) \neq \lambda_A^{\mathbf{OP}}(n)$
 - (e2) if $m \vdash_A n$ where $\lambda_A^{\mathbf{QA}}(n) = \mathbf{A}$ then $\lambda_A^{\mathbf{QA}}(m) = \mathbf{Q}$

The simplest possible arena, written $\mathbf{1}$, has the empty set as its set of moves. Given any countable set X , the **flat game** over X has one move \mathbf{q} (labelled \mathbf{OQ}) and, for each $x \in X$, a move \mathbf{x} (labelled \mathbf{PA}), the only initial token being \mathbf{q} . The enabling relation specifies $\mathbf{q} \vdash \mathbf{x}$ for all $x \in X$. We write **bool** and **nat** for the flat games over $\{\mathbf{t}, \mathbf{ff}\}$ and $\{0, 1, 2, \dots\}$ respectively.

A **legal play** of A is a pointing string s over alphabet M_A satisfying

- if s_i points to s_j then $s_j \vdash_A s_i$ and if s_i has no pointer then $s_i \in I_A$
- the underlying string of s satisfies *alternation*: if $\mathcal{U}s = s_1 m_1 m_2 s_2$ then $\lambda_A^{\mathbf{OP}}(m_1) \neq \lambda_A^{\mathbf{OP}}(m_2)$.

Each occurrence in a legal play s is an element m of M_A , with or without a pointer. If m has no pointer, we call it an **initial move** and note that $m \in I_A$; otherwise, we call m plus its pointer a **move** of s . If s_j points to s_i we sometimes say that s_i **justifies** s_j ; more generally, if following pointers back from s_j arrives at s_i , we say that s_i **hereditarily justifies** s_j . The first move of a legal play must be initial (since it cannot point to any previous move!) and hence is an \mathbf{O} -move and so alternation just means: $\lambda_A^{\mathbf{OP}}(s_i) = \mathbf{O}$ if, and only if, i is odd.

We write \mathcal{L}_A for the set of all legal plays of A . Given $s \in \mathcal{L}_A$, set $\text{ie}(s) = \{t \in \mathcal{L}_A \mid s \sqsubseteq t \wedge |t| = |s| + 1\}$, the **immediate extensions** of s . We write $\text{ip}(s)$ for the **immediate prefix** of s (provided $s \neq \varepsilon$) and (unless s is empty or ends with an initial move) $\text{jp}(s)$ for the **justifying prefix** of s , that prefix of s ending with the move that justifies s_ω . Finally, for $s, t \in \mathcal{L}_A$, we write $s \sqsubseteq^P t$ (resp. $s \sqsubseteq^O t$) iff s is a prefix of t ending with a P- (resp. O-) move. We fix the convention that $\varepsilon \sqsubseteq^P s$ for any $s \in \mathcal{L}_A$.

In game semantics, a legal play in arena A corresponds to a possible “execution trace” of a program of type A . This trace consists of alternate moves from Opponent (representing the context of execution) and Player (representing the program). We model programs as *strategies*: sets of P-ending legal plays. If s is a legal play in such a set, we interpret this to mean: if, after Opponent’s last move, the trace to date is $\text{ip}(s)$ then Player responds with s_ω (plus the appropriate pointer). So, in game semantics, we think of a program as a “lookup table” explaining how Player must respond to Opponent’s choice of moves:

A **strategy** for arena A is a non-empty set σ of P-ending legal plays of A satisfying

- if $s \in \sigma$ and $s' \sqsubseteq^{\text{even}} s$ then $s' \in \sigma$;
- if $s, t \in \sigma$ then $s \wedge t \in \sigma$.

We write $\text{dom}(\sigma)$ for $\bigcup_{s \in \sigma} \text{ie}(s)$, the *domain* of σ . The second condition says that P has *at most one response* at any given moment—what we usually call *determinism*.

Finally, we present three constructors on arenas of use later on. The **product** of arenas A and B just inherits the structure of A and B whereas the **arrow** constructor $A \Rightarrow B$ must add in some new structure to deal with its inversion of labelling on the LHS.

- $M_{A \times B} = M_A + M_B$, the disjoint union
- $\lambda_{A \times B} = [\lambda_A, \lambda_B]$, the copairing
- $I_{A \times B} = I_A + I_B$
- $\vdash_{A \times B} = \vdash_A + \vdash_B$, the sum of relations.

- $M_{A \Rightarrow B} = M_A + M_B$
- $\lambda_{A \Rightarrow B} = [\langle \lambda_A^{\text{PO}}, \lambda_A^{\text{QA}} \rangle, \lambda_B]$
- $I_{A \Rightarrow B} = I_B$
- $\vdash_{A \Rightarrow B} = (\vdash_A + \vdash_B) \cup \{(\text{inr}(b), \text{inl}(a)) \mid b \in I_B \wedge a \in I_A\}$

We also write A^- for the “decapitated” arena obtained by removing the initial moves of A , flipping the labelling function and setting the **secondary** tokens of A —all those tokens enabled by some initial token—as the initial tokens:

- $M_{A^-} = M_A - I_A$, the set difference
- $\lambda_{A^-} = \lambda_A^{\text{PO}} \upharpoonright M_{A^-}$, the restriction of λ_A^{PO} to M_{A^-}
- $I_{A^-} = \{m \in M_{A^-} \mid \exists i \in I_A. i \vdash_A m\}$
- $\vdash_{A^-} = \vdash_A \upharpoonright (M_{A^-} \times M_{A^-})$

2.2 The big SMCC of arenas and strategies

If we have strategies $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$, we can make them interact on the “common ground” of B . We formalize this idea with the notion of a **legal interaction** in arenas A , B and C : this is a pointing string u over alphabet $M_A + M_B + M_C$ such that $u \upharpoonright A, B \in \mathcal{L}_{A \Rightarrow B}$, $u \upharpoonright B, C \in \mathcal{L}_{B \Rightarrow C}$ and $u \upharpoonright A, C$ alternates (it easily follows that $u \upharpoonright A, C \in \mathcal{L}_{A \Rightarrow C}$). This definition obviously generalizes to any finite *list* of arenas. We write $\mathcal{I}(A, B, C)$ for the set of all legal interactions over A , B and C .

Intuitively, a legal interaction consists of a sequence of “sandwiches” of the form “O-move of $A \Rightarrow C$ followed by moves of B , followed by a P-move of $A \Rightarrow C$ ”. If the number of B -moves is odd, one of the outer moves comes from A and the other from C whereas, if the number of B -moves is even, both outer moves come from the “same side”. If A is the empty arena and C a flat arena, we call $u \in \mathcal{I}(A, B, C)$ a **program interaction**.

The composite of $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ is defined by setting

- $\sigma \parallel \tau = \{u \in \mathcal{I}(A, B, C) \mid u \upharpoonright A, B \in \sigma \wedge u \upharpoonright B, C \in \tau\}$
- $\sigma ; \tau = \{u \upharpoonright A, C \mid u \in \sigma \parallel \tau\}$

In words, $\sigma ; \tau$ consists of the external projections of all interactions that σ and τ mutually accept. Note that all actual interaction between σ and τ gets excised from their composite, leaving a sequence of external responses, by P in $A \Rightarrow C$, to moves of O in $A \Rightarrow C$ —the “bread” of the sandwiches.

Proposition 2.2.1 *If σ and τ are strategies for $A \Rightarrow B$ and $B \Rightarrow C$ respectively then $\sigma ; \tau$ is a strategy for $A \Rightarrow C$.*

The proof depends on the unique witness lemma [to show that determinism is preserved by composition]:

Lemma 2.2.2 (unique witness) *If σ and τ are strategies for $A \Rightarrow B$ and $B \Rightarrow C$ respectively then, for all $s \in \sigma ; \tau$, there exists a unique $u \in \mathcal{I}(A, B, C)$ such that $s = u \upharpoonright A, C$, $u \upharpoonright A, B \in \sigma$ and $u \upharpoonright B, C \in \tau$.*

We organize the above development into a category \mathbf{G} . Its objects are arenas, an arrow $f : A \rightarrow B$ is a strategy $\sigma_f : A \Rightarrow B$ and the composite of arrows $\sigma_f : A \Rightarrow B$ and $\sigma_g : B \Rightarrow C$ is $\sigma_f ; \sigma_g$. The identity for A is the copycat strategy on $A \Rightarrow A$ and associativity of composition follows from the zipping lemma:

Lemma 2.2.3 *Given $u \in \mathcal{I}(A, C, D)$ and $v \in \mathcal{I}(A, B, C)$ such that $u \upharpoonright A, C = v \upharpoonright A, C$, there exists a unique $w \in \mathcal{I}(A, B, C, D)$ such that $u = w \upharpoonright A, C, D$ and $v = w \upharpoonright A, B, C$.*

We extend the \times constructor on arenas to a bifunctor by setting, for strategies $\sigma : A \Rightarrow B$ and $\sigma' : A' \Rightarrow B'$,

$$\sigma \times \sigma' = \{s \in \mathcal{L}_{(A \times A') \Rightarrow (B \times B')} \mid s \upharpoonright A, B \in \sigma \wedge s \upharpoonright A', B' \in \sigma'\}.$$

Bifunctionality follows, in a manner analogous to associativity of composition, from a second zipping lemma.

The unit for \times is the empty arena $\mathbf{1}$ equipped with evident copycat strategies $\lambda_A : \mathbf{1} \times A \Rightarrow A$ and $\rho_A : A \times \mathbf{1} \Rightarrow A$. Since $\mathbf{1}$ is a terminal object of \mathbf{G} , we have canonical projections from $A \times B$. The associativity and commutativity isomorphisms α_{ABC} and γ_{AB} come from the natural isomorphisms $(M_A + M_B) + M_C \cong M_A + (M_B + M_C)$ and $M_A + M_B \cong M_B + M_A$ in \mathbf{Set} . Finally, since the only difference between $(A \times B) \Rightarrow C$ and $A \Rightarrow (B \Rightarrow C)$ lies in the tagging of the disjoint unions, $\mathbf{G}(A \times B, C) \cong \mathbf{G}(A, B \Rightarrow C)$, the familiar *currying* isomorphism $\Lambda(-)$. Uncurrying $\text{id}_{A \Rightarrow B}$ yields $\epsilon_{AB} : (A \Rightarrow B) \times A \Rightarrow B$, the *evaluation* map of \mathbf{G} satisfying, for any $\sigma : A \times B \Rightarrow C$, $\sigma = (\Lambda(\sigma) \times \text{id}_B); \epsilon_{BC}$.

This establishes that the category of arenas and strategies, once equipped with the above-mentioned structure, is an SMCC.

2.3 Innocent strategies

In this section, we introduce the important subclass of innocent strategies as being those strategies σ whose response to some $s \in \text{dom}(\sigma)$ depends only on a partial history of s , the so-called **P-view** $\lceil s \rceil$ of s , defined inductively:

- $\lceil s \rceil = s_\omega$, if s_ω is an initial move
- $\lceil s \rceil = \lceil \text{jp}(s) \rceil \cdot s_\omega$, if s_ω is a non-initial **O**-move
- $\lceil s \rceil = \lceil \text{ip}(s) \rceil s_\omega$, if s_ω is a **P**-move

In words, we follow pointers from **O**-moves, excising everything that lies “under the arc” of the pointer, and “step over” **P**-moves, thus tracing out a subsequence of s where **O** always points to the previous occurrence. In the last clause, s_ω points to a move in $\lceil \text{ip}(s) \rceil$ iff it happens to point to that move in s ; if s_ω points outside of $\lceil \text{ip}(s) \rceil$ in s , it doesn’t point anywhere in $\lceil s \rceil$. So the **P-view** extracts a subsequence of the underlying string which preserves pointers whenever it can.

Since pointers from **P**-moves can disappear when we calculate a **P-view**, the **P-view** of a legal play need not itself be legal. We say that a legal play s satisfies **P-visibility** iff $\lceil s \rceil$ is a legal play; so no **P**-move of s that occurs in $\lceil s \rceil$ loses its pointer in $\lceil s \rceil$. The **P-views** of arena A are those legal plays where all of **O**’s moves point to the immediately preceding move.

Given a set of coherent **P-views** V —*i.e.* where, for all $s, t \in V$, $s \wedge t$ has even length—we define a strategy $\text{tr}(V)$, all of whose plays satisfy **P-visibility**, where the response of σ to $t \in \text{dom}(\sigma)$ depends only on $\lceil t \rceil$:

- $T_0(V) = \{\varepsilon\}$
- $T_{n+1}(V) = \{s \in \mathcal{L}_A \mid \text{ip}^2(s) \in T_n(V) \wedge \lceil s \rceil \in V\}$
- $\text{tr}(V) = \bigcup_{n \in \mathbf{N}} T_n(V)$

If all plays of a strategy σ satisfy **P-visibility** then σ itself is said to satisfy **P-visibility**. Such a strategy determines a set of **P-views** $\lceil \sigma \rceil = \{\lceil s \rceil \mid s \in \sigma\}$. In general, $\lceil \sigma \rceil$ may be nondeterministic—even if σ is deterministic—but we always have that $\sigma \subseteq \text{tr}(\lceil \sigma \rceil)$.

A strategy satisfying P-visibility for which $\sigma = \text{tr}(\ulcorner\sigma\urcorner)$ is called an **innocent** strategy. In this case, $\ulcorner\sigma\urcorner$ is itself a strategy, the **view function** of σ . Innocent strategies can equivalently be characterized as those satisfying P-visibility plus

$$s \in \sigma \wedge t \in \text{dom}(\sigma) \wedge \ulcorner\text{ip}(s)\urcorner = \ulcorner t \urcorner \Rightarrow \text{match}(s, t) \in \sigma$$

where $\text{match}(s, t)$ is the unique extension of t such that $\ulcorner s \urcorner = \ulcorner \text{match}(s, t) \urcorner$; or indeed as those strategies satisfying P-visibility plus

$$s \in \mathcal{L}_A \wedge \text{ip}(s) \in \text{dom}(\sigma) \wedge \ulcorner s \urcorner \in \ulcorner \sigma \urcorner \Rightarrow s \in \sigma.$$

In words, the response of an innocent strategy depends only on the current P-view (cf. the above definition of tr). One can go on to show that tr and fn form a Galois connection: $\text{tr} \circ \text{fn}$ is a closure operator (returning the smallest innocent strategy containing its input) and $\text{fn} \circ \text{tr}$ is an interior operator (returning the largest view function contained in its input).

2.4 The CCC of innocent strategies

A move in a legal interaction $u \in \mathcal{I}(A, B, C)$ is a **generalized O-move** (resp. **generalized P-move**) in **component** \mathcal{L} (resp. \mathcal{R}) iff it is an O- (resp. P-)move of $A \Rightarrow B$ (resp. $B \Rightarrow C$). So an O-move of $A \Rightarrow C$ or any move of B is a generalized O-move (in the appropriate component) while a P-move of $A \Rightarrow C$ or any move of B is a generalized P-move (in the appropriate component). A move in B is thus a generalized O-move in one component and a generalized P-move in the other. Note that a generalized P-move is *always* immediately preceded by a generalized O-move in the *same component*. Finally, we say that moves from A and C are **external** moves of u and that moves from B are **internal** moves of u .

We extend the notion of **P-view** to legal interaction $u \in \mathcal{I}(A, B, C)$ with the following inductive definition.

- $\ulcorner u \urcorner = u_\omega$, if u_ω is an initial move of C
- $\ulcorner u \urcorner = \ulcorner \text{jp}(u) \urcorner \cdot u_\omega$, if u_ω is an external O-move of u
- $\ulcorner u \urcorner = \ulcorner \text{ip}(u) \urcorner u_\omega$, if u_ω is a generalized P-move

Recall that a legal interaction can be viewed as a sequence of sandwiches: O-move of $A \Rightarrow C$, moves of B , P-move of $A \Rightarrow C$. The P-view of a legal interaction thus consists of a subsequence of sandwiches determined by the pointers from external O-moves of u . Each sandwich of u is either removed entirely or left untouched in the P-view.

Just as the P-view of a legal play can lose pointers from P-moves and so need not itself be a legal play, $\ulcorner u \urcorner$ can lose pointers from its *generalized* P-moves and so need not be a legal interaction. However, if u results from the interaction of two P-vis strategies, *i.e.* all P-ending prefixes of *both* internal projections $u \upharpoonright A, B$ and $u \upharpoonright B, C$ satisfy P-visibility, then the following lemma guarantees that $\ulcorner u \urcorner$ is legal:

Lemma 2.4.1 *If $u \in \mathcal{I}(A, B, C)$ such that $u \upharpoonright A, B \in \sigma$ and $u \upharpoonright B, C \in \tau$ for P-vis σ and τ , and m is a generalized P-move of u in component* X , then $\ulcorner u_{\leq m} \urcorner \in \mathcal{I}(A, B, C)$.*

Proof By induction on the length of $u_{\leq m}$; two cases, depending on the move n immediately preceding m in u . If n is a generalized P-move of u (necessarily in component Y), we apply the inductive hypothesis to get $\ulcorner u_{< m} \urcorner \in \mathcal{I}(A, B, C)$. If n is an O-move of $A \Rightarrow C$, either it's initial (in which case the claim is trivial—this encompasses the base case) or it's justified by ℓ , in which case we apply the inductive hypothesis to $u_{\leq \ell}$ to get $\ulcorner u_{< m} \urcorner = \ulcorner u_{\leq \ell} \urcorner \cdot n \in \mathcal{I}(A, B, C)$. So $\ulcorner u_{< m} \urcorner \in \mathcal{I}(A, B, C)$.

By P-visibility, m points in $\ulcorner u_{< m} \upharpoonright X \urcorner$. The last move of this, n , points to a generalized P-move m' in X which occurs in $\ulcorner u_{< m} \urcorner$. In turn, m' is immediately preceded by a generalized O-move in X , n' , which also occurs in $\ulcorner u_{< m} \urcorner$. If we continue in this way, we find that $\ulcorner u_{< m} \upharpoonright X \urcorner$ is a subsequence of $\ulcorner u_{< m} \urcorner$. Hence we can attach m to the end of $\ulcorner u_{< m} \urcorner$ with correct justification pointer, yielding $\ulcorner u_{\leq m} \urcorner \in \mathcal{I}(A, B, C)$ as required. \blacksquare

Note that, for $u \in \mathcal{I}(A, B, C)$ satisfying the hypotheses of this lemma, $\ulcorner u \urcorner \upharpoonright A, C = \ulcorner u \upharpoonright A, C \urcorner$. So, as an immediate corollary of this lemma, we have that P-visibility is preserved by composition: the composite of P-vis $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ is again a P-vis strategy since, if $s \in \sigma ; \tau$, its unique witness $u \in \mathcal{I}(A, B, C)$ satisfies the hypotheses of the above lemma; hence $\ulcorner u \urcorner \in \mathcal{I}(A, B, C)$ and so $\ulcorner s \urcorner = \ulcorner u \urcorner \upharpoonright A, C \in \mathcal{L}_{A \Rightarrow C}$.

In the previous section, we gave a way of defining innocent strategies as those (necessarily) P-vis strategies entirely determined by their view function—by “completing” the P-vis strategy with the closure operator $\text{tr} \circ \text{fn}$. We also noted an equivalent characterization: a P-vis strategy σ is innocent iff

$$s \in \sigma \wedge t \in \text{dom}(\sigma) \wedge \ulcorner \text{ip}(s) \urcorner = \ulcorner t \urcorner \Rightarrow \text{match}(s, t) \in \sigma$$

In this definition, the play $s \in \sigma$ could be any legal play but in fact it suffices to consider P-views (since $\text{match}(s, t) = \text{match}(\ulcorner s \urcorner, t)$). This observation leads to a third, perhaps simpler, characterization of innocence as P-visibility plus

$$s \in \mathcal{L}_A \wedge \text{ip}(s) \in \text{dom}(\sigma) \wedge \ulcorner s \urcorner \in \ulcorner \sigma \urcorner \Rightarrow s \in \sigma.$$

To see the equivalence of these two definitions, let σ be a P-vis strategy for A . If σ satisfies the first definition, suppose we have some $s \in \mathcal{L}_A$ such that $\text{ip}(s) \in \text{dom}(\sigma)$ and $\ulcorner s \urcorner \in \ulcorner \sigma \urcorner$. So, for some $t \in \sigma$, $\ulcorner t \urcorner = \ulcorner s \urcorner$ and hence $s = \text{match}(t, \text{ip}(s)) \in \sigma$ as required. For the other direction, if $s \in \sigma$, $t \in \text{dom}(\sigma)$ where $\ulcorner \text{ip}(s) \urcorner = \ulcorner t \urcorner$ then $\text{match}(s, t) \in \mathcal{L}_A$, $\text{ip}(\text{match}(s, t)) = t \in \text{dom}(\sigma)$ and $\ulcorner \text{match}(s, t) \urcorner = \ulcorner s \urcorner \in \ulcorner \sigma \urcorner$. Hence $\text{match}(s, t) \in \sigma$ as required.

* We use X as a metavariable ranging over $\{\mathcal{L}, \mathcal{R}\}$ when we neither know nor care which of the two components is being referred to, whence we use Y to denote *the other* component.

The next lemma plays a vital role in showing that innocent strategies compose.

Lemma 2.4.2 *If $u \in \mathcal{I}(A, B, C)$ such that $u \upharpoonright A, B \in \sigma$ and $u \upharpoonright B, C \in \tau$ for P -vis σ and τ , and m is a generalized O -move of u in component X , then $\ulcorner u_{\leq m} \upharpoonright X \urcorner = \ulcorner u_{\leq m} \urcorner \upharpoonright X \urcorner$.*

Proof By induction on the length of $u_{\leq m}$. If m is an initial move of $u \upharpoonright X$ then $\ulcorner u_{\leq m} \upharpoonright X \urcorner = m = \ulcorner u_{\leq m} \urcorner \upharpoonright X \urcorner$. Otherwise, m is either an O -move of $A \Rightarrow C$ or a generalized P -move in component Y . In either case, the justifier n of m occurs in $\ulcorner u_{\leq m} \urcorner$.

We finish by calculating

$$\begin{aligned} \ulcorner u_{\leq m} \urcorner \upharpoonright X \urcorner &= \ulcorner u_{\leq n} \urcorner \upharpoonright X \urcorner \cdot m \\ &= \ulcorner u_{< n} \urcorner \upharpoonright X \urcorner n \cdot m \\ &= \ulcorner u_{< n} \upharpoonright X \urcorner n \cdot m \\ &= \ulcorner u_{\leq n} \upharpoonright X \urcorner \cdot m \\ &= \ulcorner u_{\leq m} \upharpoonright X \urcorner \end{aligned}$$

using the inductive hypothesis and the definition of P -view. ■

Essentially, this lemma says that $\ulcorner u \urcorner$ not only witnesses $\ulcorner u \upharpoonright A, C \urcorner$ but that it contains all the information needed for two interacting *innocent strategies* to build a witness for $\ulcorner u \upharpoonright A, C \urcorner$.

This allows us to extend the function `match` to legal interactions. Suppose we have $u, v \in \sigma \parallel \tau \subseteq \mathcal{I}(A, B, C)$ where σ and τ are innocent strategies, u witnesses $s \in \sigma ; \tau$ and v is the minimum witness of $t \in \mathbf{dom}(\sigma ; \tau)$ —a legal interaction ending with an external O -move. Suppose further that $\ulcorner v \urcorner = \ulcorner \mathbf{trnct}(u) \urcorner$ where $\mathbf{trnct}(u)$ is that legal interaction obtained by removing the “tail” of generalized P -moves from u , *i.e.* the minimum witness of $\mathbf{ip}(s)$.

Let X be the component where the last move of v occurs. By the above lemma, we know that $\ulcorner v \upharpoonright X \urcorner = \ulcorner v \urcorner \upharpoonright X \urcorner = \ulcorner \mathbf{trnct}(u) \urcorner \upharpoonright X \urcorner = \ulcorner u \upharpoonright X \urcorner$. Since σ and τ are innocent, the response (in component X) to $\ulcorner v \upharpoonright X \urcorner$ is the same as to $\ulcorner u \upharpoonright X \urcorner$. This reasoning can be iterated, alternating between the two components, to extend v with the same “tail” of generalized P -moves that was removed in $\mathbf{trnct}(u)$. We thus end up with a legal interaction `match`(u, v) witnessing that $t' \in \mathbf{ie}(t)$ such that $\ulcorner t' \urcorner = \ulcorner s \urcorner$.

Proposition 2.4.3 *If $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ are innocent strategies then so is their composite $\sigma ; \tau : A \Rightarrow C$.*

Proof We use the third (and final) characterization of innocence. Suppose $s \in \mathcal{L}_{A \Rightarrow C}$ where $\mathbf{ip}(s) \in \mathbf{dom}(\sigma ; \tau)$ and $\ulcorner s \urcorner \in \ulcorner \sigma ; \tau \urcorner$. Then we have some $t \in \sigma ; \tau$, witnessed by $v \in \sigma \parallel \tau$, such that $\ulcorner t \urcorner = \ulcorner s \urcorner$. Moreover, we have a minimum witness u for $\mathbf{ip}(s)$.

So $\ulcorner u \urcorner \upharpoonright A, C = \ulcorner \mathbf{ip}(s) \urcorner = \ulcorner \mathbf{ip}(t) \urcorner = \ulcorner \mathbf{trnct}(v) \urcorner \upharpoonright A, C$ and, by the unique witness lemma 2.2.2, we conclude that $\ulcorner u \urcorner = \ulcorner \mathbf{trnct}(v) \urcorner$. Hence `match`(v, u) witnesses `match`($t, \mathbf{ip}(s)$) = $s \in \sigma ; \tau$. ■

The resulting category \mathbf{I} of arenas and innocent strategies is Cartesian closed for general reasons [3] but, more concretely, it can readily be verified that $\langle \sigma, \tau \rangle = \{s \in \mathcal{L}_{A \Rightarrow (B \times C)} \mid s \upharpoonright A, B \in \sigma \wedge s \upharpoonright A, C \in \tau\}$ defines the appropriately unique pairing of $\sigma : A \Rightarrow B$ and $\tau : A \Rightarrow C$. The **diagonal** transformation of \mathbf{G} , with components $\Delta_A = \langle \text{id}_A, \text{id}_A \rangle$, is a *natural* transformation in \mathbf{I} .

2.5 The visibility condition

The **long O-view** of $s \in \mathcal{L}_A$, noted $\lfloor s \rfloor$, is defined dually to the P-view: we follow pointers from P-moves so this time O-moves can lose pointers:

- $\lfloor \varepsilon \rfloor = \varepsilon$
- $\lfloor s \rfloor = \lfloor \text{jp}(s) \rfloor \cdot s_\omega$, if s_ω is a P-move
- $\lfloor s \rfloor = \lfloor \text{ip}(s) \rfloor s_\omega$, if s_ω is an O-move

In contrast to a P-view which always has a unique initial move, a long O-view may contain many initial moves. A play $s \in \mathcal{L}_A$ satisfies **O-visibility** iff $\lfloor s \rfloor \in \mathcal{L}_A$ —so we lose no O-pointers in $\lfloor s \rfloor$ —and satisfies simply **visibility** iff it satisfies P- and O-visibility.

Lemma 2.5.1 *Let $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ be P-vis strategies and suppose that $u \in \sigma \parallel \tau$ such that, for all external O-moves o of u , we have that $u_{\leq o} \upharpoonright A, C$ satisfies O-visibility. Then, for a non-initial external move m in component X , $\lfloor u_{\leq m} \upharpoonright Y \rfloor$ extends $\lfloor \text{jp}(u_{\leq m}) \upharpoonright Y \rfloor$.*

Proof By induction on the length of $u_{\leq m}$ (for m an external move). The base case for m a P-move consists of a single sandwich which has the required property thanks to P-visibility in X . The base case for m an O-move consists of a sandwich plus a single external O-move which trivially has the required property.

If m is an external P-move in component X , its justifier occurs in $\lceil u_{< m} \upharpoonright X \rceil$. If we trace back $\lfloor u_{< m} \upharpoonright Y \rfloor$, it starts out just like $\lceil u_{< m} \upharpoonright X \rceil$ until this latter steps back onto an *external* O-move o in X . At this point, we can apply the inductive hypothesis to $u_{\leq o}$ so that $\lfloor u_{\leq o} \upharpoonright Y \rfloor$ extends $\lfloor \text{jp}(u_{\leq o}) \upharpoonright Y \rfloor$. (In other words, following back the O-view of $u_{\leq o} \upharpoonright Y$ can never jump past the justifier of o .) If we continue in this way, we eventually arrive at the last move of $\text{jp}(u_{\leq m}) \upharpoonright Y$ and so $\lfloor u_{\leq m} \upharpoonright Y \rfloor$ extends $\lfloor \text{jp}(u_{\leq m}) \upharpoonright Y \rfloor$ as required.

If m is an external O-move in component X , external O-visibility implies that its justifier lies in $\lfloor u_{< m} \upharpoonright A, C \rfloor$. As we trace back $\lfloor u_{\leq m} \upharpoonright Y \rfloor$, we apply the inductive hypothesis to $u_{\leq p}$ for all the external P-moves p of $\lfloor u_{< m} \upharpoonright A, C \rfloor$ (in component X) that lie between m and its justifier. This establishes, at each such point, that $\lfloor u_{\leq p} \upharpoonright Y \rfloor$ extends $\lfloor \text{jp}(u_{\leq p}) \upharpoonright Y \rfloor$. Hence $\lfloor u_{\leq m} \upharpoonright Y \rfloor$ extends $\lfloor \text{jp}(u_{\leq m}) \upharpoonright Y \rfloor$ as required. \blacksquare

This lemma basically says that the pointers from *external* moves in component X constrain the pointers in component Y , even though external moves in X cannot be seen from Y !

In an innocent strategy, all plays satisfy P-visibility but not necessarily O-visibility. In other words, an innocent strategy does not need its *context* to respect P-visibility. However, in order to build a category of innocent, or just P-vis, strategies, we have no need for this extra generality: when P-vis σ and τ interact, σ appears O-vis from τ 's point of view and τ appears O-vis from σ 's point of view, so we can safely ignore all plays violating O-visibility:

Lemma 2.5.2 *Let $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ be P-vis strategies and suppose that $u \in \sigma \parallel \tau$ such that, for all external O-moves o of u , we have that $u_{\leq o} \upharpoonright A, C$ satisfies O-visibility. Then, for any generalized O-move m of u in component X , we have that $u_{\leq m} \upharpoonright X$ satisfies O-visibility.*

Proof If m is an external move, it's either initial (whence we have nothing to prove) or it points in $[u_{< m} \upharpoonright A, C]$ in component X . For each P-move of $[u_{< m} \upharpoonright A, C]$ in component Y , we apply the above lemma, thus establishing that m points in $[u_{< m} \upharpoonright X]$ as required.

Otherwise, m is a generalized O-move in X which is also a generalized P-move in Y and so, by P-visibility, points in $\lceil u_{< m} \upharpoonright Y \rceil$. We apply the above lemma to each external O-move of $\lceil u_{< m} \upharpoonright Y \rceil$, thus establishing that m points in $[u_{< m} \upharpoonright X]$ as required. ■

This implies, for innocent (or just P-vis) strategies $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$, that any $s \in \sigma ; \tau$ that happens to satisfy O-visibility necessarily comes from an interaction between O-vis plays of σ and τ . This enables us [G. McCusker, private communication] to define a category of P- and O-vis strategies which, qua category, differs inessentially from the category of P-vis strategies.


2.6 The annotated O-view

Earlier on, we defined the function $\text{match}(s, t)$ to formalize what we intuitively mean by “extend t with the last move of s ”. We now define a new function $\text{match}^*(s, t)$ for P-ending, P-vis $s \in \mathcal{L}_A$ and $t \in \mathcal{L}_A$ such that $\lceil \text{jp}(s) \rceil = \lceil t \rceil$: $\text{match}^*(s, t)$ denotes the extension of t with that suffix of the P-view of s that lies underneath the pointer from s_ω to $\text{jp}(s)_\omega$. In other words, instead of adding just the last move of s to t , this adds the last **arch** of the P-view of s to t .

The **annotated (long) O-view** of P- and O-vis $s \in \mathcal{L}_A$, written $\widetilde{[s]}$, is defined inductively:

- $\widetilde{[\varepsilon]} = \varepsilon$
- $\widetilde{[s]} = \widetilde{[\text{ip}(s)]}_{s_\omega}$, if s_ω is an O-move
- $\widetilde{[s]} = \text{match}^*(s, \widetilde{[\text{jp}(s)]})$, if s_ω is a P-move

In words, the annotated (long) O-view traces back the (long) O-view but, instead of *excising* all moves underneath each P-to-O pointer, it retains that suffix of the current P-view “enclosed” by the pointer—the last arch:



The assumption of \mathbf{O} -visibility for s implies that $\widetilde{[s]} \in \mathcal{L}_A$. So $\widetilde{[s]}$ in fact consists of an *interleaving of P-views* of A . If $s \in \sigma$ for an innocent σ then $\widetilde{[s]}$ exposes what input P-view σ requires in order to produce $[s]$ as output \mathbf{O} -view.

We previously defined $\text{match}(u, v)$ to formalize what we mean by “extend v with the tail of generalized P-moves of u ”. In order to define $\text{match}^*(u, v)$, consider $u \in \sigma \parallel \tau$ (for P-vis $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$) and $v \in \mathcal{I}(A, B, C)$ such that $\lceil \text{jp}(u) \rceil = \lceil v \rceil$. We define $\text{match}^*(u, v)$ to be the extension of v with the last arch of the $\lceil u \rceil$, *i.e.* extend v with the suffix of $\lceil u \rceil$ lying underneath the pointer from u_ω .

The P-view of a legal interaction picks out a subsequence of sandwiches according to the following scheme: track back through the current sandwich until we reach an external \mathbf{O} -move o ; then follow o 's pointer and recursively apply, until we reach an initial \mathbf{O} -move. The (long) \mathbf{O} -view [defined below] of a legal interaction u picks out its subsequence of sandwiches differently: it applies the above procedure (for the P-view) but only until we reach that external \mathbf{O} -move o that justifies the last move of u —assumed to be an external P-move; then move to the external P-move immediately preceding o and recursively apply.

Suppose that our $u \in \sigma \parallel \tau$ additionally satisfies external \mathbf{O} -visibility: for all external \mathbf{O} -moves o of u , we have $u_{\leq o} \upharpoonright A, C \in \mathcal{L}_{A \Rightarrow C}$. We define the (**long**) \mathbf{O} -**view** of u , written $[u]$, inductively:

- $[\varepsilon] = \varepsilon$
- $[u] = [\text{ip}(u)]u_\omega$, if u_ω is an external \mathbf{O} -move
- $[u] = \text{match}^*(u, [\text{jp}(u)])$, if u_ω is an external P-move

Note that $[u]$ is an interleaving of P-views of legal interactions just as $\widetilde{[s]}$ is an interleaving of P-views of legal plays. External \mathbf{O} -visibility guarantees that $[u] \in \mathcal{I}(A, B, C)$. Moreover, $\lceil u \rceil$ is a subsequence of $[u]$ and $[u]$ clearly witnesses $[u \upharpoonright A, C]$.

2.7 Single-threaded and interference-free strategies

The theorem that establishes that innocent strategies form a CCC applies more generally to *single-threaded* strategies [1, 4]. These lie in one-two-one correspondence with the **comonoïd homomorphisms**, *i.e.* arrows $\sigma : A \rightarrow B$ satisfying

$$\begin{array}{ccc} A & \xrightarrow{\Delta_A} & A \times A \\ \sigma \downarrow & & \downarrow \sigma \times \sigma \\ B & \xrightarrow{\Delta_B} & B \times B \end{array} \qquad \begin{array}{ccc} A & & \\ \sigma \downarrow & \searrow !_A & \\ B & \xrightarrow{!_B} & \mathbf{1} \end{array}$$

where $\Delta_A : A \rightarrow A \times A$ is (a component of) the diagonal transform. (In fact, since $\mathbf{1}$ is terminal, the triangle always commutes.)

Briefly, a **thread** of an arena A is a legal play of A with (at most) one initial move. We write $[\sigma]$ for the set of threads contained in strategy σ . Given a coherent set of even-length threads T , define

- $\text{ST}_0(T) = \{\varepsilon\}$
- $\text{ST}_{i+1}(T) = \{s \in \mathcal{L}_A \mid \text{ip}^2(s) \in \text{ST}_i(T) \wedge [s] \in T\}$
- $\text{ST}(T) = \bigcup_{i \in \mathbb{N}} \text{ST}_i(T)$.

In a manner analogous to the characterization of innocent strategies, we say that σ is **single-threaded** if, and only if, $\sigma = \text{ST}[\sigma]$. We refer to $[\sigma]$ as the **thread function** of σ (analogous to the *view function* of an innocent strategy).

If σ and τ are composable single-threaded strategies, $\sigma; \tau$ is also single-threaded (comonoid homomorphisms can easily be seen to be closed under composition) with thread function $[\sigma; \tau] = \sigma; [\tau]$. Note that a single thread of $\sigma; \tau$ may depend on an arbitrary play, not necessarily just a thread, of σ .

We say that a legal play is **well-opened** iff no initial move is ever repeated, *i.e.* distinct initial moves are occurrences of different initial tokens. Beware that this definition conflicts with much of the literature of game semantics which conflates the concepts of thread and well-opened play.

We write $\mathcal{L}_A^{\text{wo}}$ for the set of all well-opened plays of A and $\text{WO}(\sigma)$ for the well-opened plays of σ . The well-opened plays of a single-threaded strategy σ thus give all possible “resource sensitive” interleavings of threads of σ : we can arbitrarily interleave threads provided their initial moves are different. So, in composition, we have $\text{WO}(\sigma; \tau) = \sigma; \text{WO}(\tau)$ where, in general, we need a non-well-opened play of σ .

If single-threaded σ has the property that, in no well-opened $s \in \sigma$ do we ever repeat a *secondary* move, then we say that σ is an **interference-free** strategy. So the *threads* of $\sigma; \tau$ are determined by interactions between well-opened plays of σ and threads of τ , the idea being that τ never opens more than one copy of any given thread of σ , a “no nesting” constraint. However, this doesn’t suffice to explain the interference-free condition which further requires that any *well-opened interleaving* of threads of τ must appear well-opened from σ ’s *point of view*, a typical non-example being Δ_A which repeats initial moves on the LHS. So, for composable interference-free σ and τ , we have

$$\text{WO}(\sigma; \tau) = \text{WO}(\sigma); \text{WO}(\tau).$$

A simple calculation establishes that interference-free strategies are closed under composition. The resulting category is an SMCC but fails to be a CCC: unlike in the category \mathbf{G} , where the diagonal transform exists but isn’t natural, it doesn’t even exist in the category of interference-free strategies.

3 Affine strategies

In the general setting of single-threaded strategies, the notion of “interference-free” isolates a class of “affine” strategies which correspond fairly closely to basic SCI (provided we also ask for P-visibility). In this section, we turn to the more specific case of innocent strategies and how to characterize those definable in (some kind of simply-typed) affine λ -calculus.

3.1 The short O-view

When two innocent strategies interact, each strategy receives input from the other, computes its view and makes its response, which it promptly sends back to the other who, in turn, does the same thing. We consider an innocent strategy σ to be “affine” iff, during this interaction, for as long as σ ’s context τ never repeats an input view for σ , σ never repeats an input view for τ .

In order to formalize this idea of “repeating a view for the other strategy”, we introduce a variant notion of O-view, in some sense dual to that of P-view. The **short O-view** $\lfloor s \rfloor$ of a non-empty $s \in \mathcal{L}_A$ is defined by:

- $\lfloor s \rfloor = \varepsilon$, if s_ω is an initial move
- $\lfloor s \rfloor = \lfloor \text{jp}(s) \rfloor \cdot s_\omega$, if s_ω is a P-move
- $\lfloor s \rfloor = \lfloor \text{ip}(s) \rfloor s_\omega$, if s_ω is a non-initial O-move

Note that tracing back $\lfloor s \rfloor$ always ends with a secondary move—so a short O-view contains a unique secondary move, just as a P-view contains a unique initial move. However, this means that $\lfloor s \rfloor$ can never be a legal play of A but, if $\lfloor s \rfloor$ doesn’t lose the pointer from any of its non-initial O-moves, then $\lfloor s \rfloor$ is a legal play of A^- . In this case, if s was played by some strategy σ for A , $\lfloor s \rfloor$ determines the next input P-view for σ ’s context.

3.2 The SMCC of affine strategies


As mentioned above, an affine strategy never repeats an input view for its context for as long as its context does the same. We formalize this idea as *injectivity of use of the view function* within a well-opened play: an innocent strategy σ for A is **affine** iff

$$s, t \in \sigma \cap \mathcal{L}_A^{\text{wo}} \wedge s \sqsubseteq t \wedge \lfloor s \rfloor = \lfloor t \rfloor \in \mathcal{L}_{A^-} \implies \lceil \text{ip}(s) \rceil = \lceil \text{ip}(t) \rceil.$$

So, if σ *does* repeat an input view for its context, *i.e.* it plays t with same O-view as an earlier s , this means that its context just did the same thing, *i.e.* $\lceil \text{ip}(s) \rceil = \lceil \text{ip}(t) \rceil$. The requirement that $\lfloor t \rfloor$ be legal in A^- means that we never lose pointers from t ’s O-occurrences while tracing back its short O-view.

Proposition 3.2.1 *If $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ are affine strategies then so is $\sigma ; \tau : A \Rightarrow C$.*

Proof Suppose $s, t \in (\sigma; \tau) \cap \mathcal{L}_{A \Rightarrow C}^{\text{wo}}$ satisfying $s \sqsubseteq t$ and $\lfloor s \rfloor = \lfloor t \rfloor$. By the definition of composition, we have witnesses $u, v \in \sigma \parallel \tau$ for s and t . Assuming that $s \neq t$ (otherwise the claim is trivial), the fact that $\lfloor s \rfloor = \lfloor t \rfloor$ implies that some P-move \bullet of $\lfloor t \rfloor$ rejoins $\lfloor s \rfloor$:




Now, $\lfloor u \rfloor \in \sigma \parallel \tau$, since σ and τ are both innocent. Note that $\lfloor u \rfloor \upharpoonright A, C \rfloor = \lfloor \lfloor u \rfloor \rfloor \upharpoonright A, C \rfloor$. We next perform the following surgery on v to obtain $\tilde{v} \in \sigma \parallel \tau$ satisfying $\lfloor u \rfloor \sqsubseteq \tilde{v}$ and $\lfloor v \rfloor \upharpoonright A, C \rfloor = \lfloor \tilde{v} \rfloor \upharpoonright A, C \rfloor$:

- follow back $\lfloor v \rfloor$ until we reach the external P-move \bullet that rejoins $\lfloor s \rfloor$
- follow back $\lceil v_{\leq \bullet} \rceil$ until we reach an external O-move \circ that points in s (and so in $\lfloor u \rfloor$, by P-visibility at \bullet_t)
- delete the remaining moves strictly inbetween the end of u and \circ and prefix the resulting sequence of moves with $\lfloor u \rfloor$; call this \tilde{v} .

By O- and P-visibility, all moves of \tilde{v} in the suffix of $\lfloor v \rfloor$ point either within $\lfloor v \rfloor$ or in $\lfloor u \rfloor$. Similarly, all moves of \tilde{v} in the suffix of $\lceil v_{\leq \bullet} \rceil$ point either in $\lceil v_{\leq \bullet} \rceil$ or in $\lfloor u \rfloor$. So $\tilde{v} \in \mathcal{I}(A, B, C)$. Moreover, by innocence of σ and τ , $\tilde{v} \in \sigma \parallel \tau$.


Consider $\lfloor u \rfloor_{\leq \bullet_s}$ and $\tilde{v}_{\leq \bullet_t}$ and suppose that the two \bullet s occur in component X . So $\lfloor \lfloor u \rfloor_{\leq \bullet_s} \rfloor \upharpoonright X \rfloor = \lfloor \tilde{v}_{\leq \bullet_t} \rfloor \upharpoonright X \rfloor$. By affinity in X , $\lceil \lfloor u \rfloor_{\leq \bullet_s} \rceil \upharpoonright X \rfloor = \lceil \tilde{v}_{\leq \bullet_t} \rceil \upharpoonright X \rfloor$ and so $\lfloor \lfloor u \rfloor_{\leq \bullet_s} \rfloor \upharpoonright Y \rfloor = \lfloor \tilde{v}_{\leq \bullet_t} \rfloor \upharpoonright Y \rfloor$. Carrying on in this way, we establish that the last arch of $\lfloor u \rfloor_{\leq \bullet_s}$ is exactly the same as the last arch of $\tilde{v}_{\leq \bullet_t}$ and so $\lfloor u \rfloor_{\leq \bullet_s} = \lfloor \tilde{v}_{\leq \bullet_t} \rfloor$. By repeatedly applying innocence (and determinism), we find that the last arches of $\lfloor u \rfloor$ are the exactly the same as those of $\lfloor \tilde{v} \rfloor$ and indeed of $\lfloor v \rfloor$ and so $\lfloor u \rfloor = \lfloor \tilde{v} \rfloor = \lfloor v \rfloor$. Hence $\lceil \text{ip}(s) \rceil = \lceil \text{ip}(t) \rceil$ as required. ■

In essence, this proof says that, if P rejoins at some point, *i.e.* $\lfloor t \rfloor$ rejoins $\lfloor s \rfloor$, then $\lceil \text{ip}(t) \rceil$ rejoins $\lceil \text{ip}(s) \rceil$, *i.e.* O rejoins at some previous point.




In a legal interaction, this property is iterated in order to establish that any repeated external O-view ultimately stems from a repeated P-view by the *external* Opponent.

We can thus form a category of arenas and affine strategies. This category can be equipped with the SMCC structure of \mathbf{G} but the diagonal transform, as in the case of interference-free strategies, no longer exists: a typical well-opened play of Δ_A repeats the trivial O-view consisting of just a secondary move:



3.3 The μ -transform

If B_1 and B_2 are flat arenas, we have an important affine strategy $\mu_{B_1 B_2}$ for $(B_1 \Rightarrow B_2) \Rightarrow (B_1 \Rightarrow B_2)$ that transforms an *innocent* strategy into an *affine* strategy by “compressing” all accesses to B_1 into one single “shared” access:



If we have an innocent strategy $\sigma : \llbracket \Gamma \rrbracket \Rightarrow (B_1 \Rightarrow B_2)$, it may interrogate its argument B_1 multiple times. If we compose σ with $\mu_{B_1 B_2}$, this has the effect of “linearizing” σ in B_1 : $\mu_{B_1 B_2}$ acts as a kind of “delayed copycat” that first evaluates the argument in B_1 “once and for all” and then, each time that σ subsequently requests that argument, $\mu_{B_1 B_2}$ simply repeats the value *already computed*. In this way, the composite $\sigma ; \mu_{B_1 B_2}$ satisfies affinity.

3.4 Affine PCF

The SMCC of affine strategies clearly harbours a model of affine λ -calculus extended with constants and special forms for arithmetic. We content ourselves with the simple type hierarchy starting from base types $B ::= \text{bool} \mid \text{nat}$.

affine lambda		
$\frac{}{\Gamma, x : T \vdash x : T}$	$\frac{\Gamma, x : T \vdash M : U}{\Gamma \vdash \lambda x(M) : T \multimap U}$	
$\frac{\Gamma \vdash M : T \multimap U \quad \Delta \vdash N : T}{\Gamma, \Delta \vdash (M)N : U}$		
arithmetic		
$\frac{}{\Gamma \vdash \text{t} : \text{bool}}$	$\frac{}{\Gamma \vdash \text{ff} : \text{bool}}$	$\frac{}{\Gamma \vdash n : \text{nat}} \quad n \in \mathbf{N}$
$\frac{\Gamma \vdash M : \text{nat}}{\Gamma \vdash (\text{succ } M) : \text{nat}}$	$\frac{\Gamma \vdash M : \text{nat} \quad \Delta \vdash N : \text{nat}}{\Gamma, \Delta \vdash (\text{equal? } M N) : \text{bool}}$	

We use `equal?` rather than `pred` and `zero?` to emphasize the “multiplicative” format of binary rules.

3.4.1 Adding conditionals

In order to establish a definability result for [compact, well-bracketed] affine strategies, we need to add special forms for conditionals (of types `bool` and `nat`). However, the usual kind of typing rule

$$\frac{\Gamma \vdash M : \text{bool} \quad \Gamma \vdash N : B \quad \Gamma \vdash L : B}{(\text{if } M \text{ } N \text{ else } L) : B}$$

becomes problematic in an affine setting: such programming constructs first evaluate their “guard” and then evaluate a “continuation” chosen on the basis of the value of the guard. Even if the guard M and continuations N and L to an `if` are all affine, the fact that they share (in the above rule) the same free variables means that violations of affinity in $(\text{if } M \text{ } N \text{ else } L)$ cannot be ruled out. This suggests use of a “multiplicative” version of the typing rule, although N and L can obviously safely share free variables since at most one of them will be evaluated:

$$\frac{\Gamma \vdash M : \text{bool} \quad \Delta \vdash N : B \quad \Delta \vdash L : B}{\Gamma, \Delta \vdash (\text{if } M \text{ } N \text{ else } L) : B}$$

Unfortunately, this leads to a loss of expressive power: if the value of the guard depends on (some of) its free variables in Γ , the selected continuation will be unable to *recompute* that value. This suggests that the *value* of the guard be passed as an extra (CBV-style) parameter to the chosen continuation:

<p>if/case</p> $\frac{\Gamma \vdash M : \text{bool} \quad \Delta, z : \text{bool} \vdash N : B \quad \Delta, z : \text{bool} \vdash L : B}{\Gamma, \Delta \vdash (\text{if } [z = M] \text{ } N \text{ else } L)}$ $\frac{\Gamma \vdash M : \text{nat} \quad \Delta, z : \text{nat} \vdash \vec{N} : B \quad \Delta, z : \text{nat} \vdash L : B}{\Gamma, \Delta \vdash (\text{case } [z = M] \vec{N} \text{ else } L) : B}$

This allows us to rewrite a typical *non-affine* PCF term

$$f : \text{nat} \rightarrow \text{nat} \vdash \lambda n (\text{case } n \text{ } 0 \text{ else } (\text{succ } (\text{succ } (f)(\text{pred } n)))) : \text{nat}$$

to evaluate n “once and for all” and bind its value to z for the benefit of the `else` continuation:

$$f : \text{nat} \rightarrow \text{nat} \vdash \lambda n (\text{case } [z = n] \text{ } 0 \text{ else } (\text{succ } (\text{succ } (f)(\text{pred } z)))) : \text{nat}$$

To clarify the difference between these terms, consider the big-step operational semantics for `if` and `case`:


$$\boxed{
\begin{array}{c}
\frac{M \Downarrow \mathbf{t} \quad N[\mathbf{t}/z] \Downarrow V}{(\mathbf{if} [z = M] N \mathbf{else} L) \Downarrow V} \quad \frac{M \Downarrow \mathbf{ff} \quad L[\mathbf{ff}/z] \Downarrow V}{(\mathbf{if} [z = M] N \mathbf{else} L) \Downarrow V} \\
\\
\frac{M \Downarrow \mathbf{m} \quad N_m[\mathbf{m}/z] \Downarrow V}{(\mathbf{case} [z = M] N_0 \cdots N_k \mathbf{else} L) \Downarrow V} \quad 0 \leq m \leq k \\
\\
\frac{M \Downarrow \mathbf{m} \quad L[\mathbf{m}/z] \Downarrow V}{(\mathbf{case} [z = M] N_0 \cdots N_k \mathbf{else} L) \Downarrow V} \quad k < m
\end{array}
}$$

Here we can clearly see the role of the free variable z in each continuation: it reifies into our language the value (\mathbf{t} , \mathbf{ff} or some \mathbf{n}) that provoked evaluation of the continuation in question. In the absence of this, a continuation has no knowledge of *why* it was selected.

On the semantic side, we must change the interpretation of \mathbf{if} (and \mathbf{case}), using the μ -transform to implement this “once and for all” evaluation of the guard. If $\Gamma \vdash M : \mathbf{bool}$ and $\Delta, z : \mathbf{bool} \vdash N : B$ and $\Delta, z : \mathbf{bool} \vdash L : B$, define $\llbracket (\mathbf{if} [z = M] N \mathbf{else} L) \rrbracket$ by composing


$$\llbracket M \rrbracket \times \langle \Lambda(\llbracket N \rrbracket) ; \mu_{\mathbf{bool}, B}, \Lambda(\llbracket L \rrbracket) ; \mu_{\mathbf{bool}, B} \rangle$$

with the following strategy for $(\mathbf{bool} \times (\mathbf{bool} \Rightarrow B) \times (\mathbf{bool} \Rightarrow B)) \Rightarrow B$:




Note how we apply the μ -transform to our continuations N and L . This has the effect of forcing the chosen continuation to immediately ask for the value of the guard that caused its evaluation.

Returning to our example, the non-affine PCF term *reevaluates* n whenever its **else** clause is selected—and yet the value of n (*i.e.* Sx) lies in the view!



With the new definition of **case**, we don't have this problem; the variable z reifies Sx so that the reference to **(pred z)** returns the value that x stands for.



In a *uniform* language like PCF or μ PCF, it obviously makes no difference for *correctness* whether we reevaluate or share. So, in PCF, the μ -transform improves efficiency without affecting expressive power whereas, in affine PCF, we absolutely need *some way* to express sharing. We choose to do this with the **if** and **case** special forms.

We could do this differently. For example, we could type terms of affine PCF in split contexts $\Gamma \mid \Delta$ where Δ is a context (containing only variable-base type pairs) of “sharable variables”. The μ -transform would then appear explicitly as the means of migrating $x : B \in \Gamma$ across to Δ , *i.e.* the syntax for saying “evaluate this input and save its value for later”, while **if** can be typed “multiplicatively” for Γ and “additively” for Δ . A term (with empty context Γ) would then be evaluated in an “environment”, *i.e.* an assignment of *values* to its *sharable* variables.

$$\frac{\Gamma, x : B \mid \Delta \vdash M : T}{\Gamma \mid \Delta, x : B \vdash \mu x(M) : T}$$

$$\frac{\Gamma \mid \Delta \vdash M : \text{bool} \quad \Gamma' \mid \Delta \vdash N : B \quad \Gamma' \mid \Delta \vdash L : B}{\Gamma, \Gamma' \mid \Delta \vdash (\text{if } M \text{ } N \text{ else } L) : B}$$

This solution, while more general than adopting the variant syntax for **if** and **case**, seems like overkill in the case of affine PCF. However, in the case of general innocent strategies, the idea of making environments into first-class objects seems rather natural as a way of formalizing *when* a given subterm gets evaluated.

Theorem 3.4.1 *If σ is a compact, affine, well-bracketed strategy on an arena of the form $(T_1 \times \cdots \times T_n) \Rightarrow T$, where T and the T_i s interpret types of affine PCF, then there exists a term M of affine PCF such that $\sigma = \llbracket M \rrbracket$.*

Proof We follow the usual pattern of such definability results [2], the key point being that, when we extract the argument and continuation strategies $\mathbf{args}(\sigma)$ and $\mathbf{cnts}(\sigma)$, affinity of σ guarantees that we can partition the context Γ of M into Γ_1 and Γ_2 where $\mathbf{args}(\sigma)$ uses only Γ_1 and $\mathbf{cnts}(\sigma)$ only uses Γ_2 . (This partition need not be unique since some variables of Γ might not be used at all.) We can then “put σ back together again” using **if** or **case** as appropriate. ■

Let us note a subtle point about the syntax of **if** and **case**: although we allow *all* continuations to use a variable to reify the value of the guard, we only ever really need this in the **else** continuation; the other continuations effectively have this value hardwired.

In the above proof, we assume a compact σ and hence the **else** continuation is always divergent and so we never need to explicitly reify the guard as a fresh variable added to Γ_2 . In other words, the naïve affine typing of **if** and **case** works perfectly adequately for “compact” terms; but once we have a non-trivial **else** clause, as in the above example, we no longer have the guard’s value hardwired and so we really need it to be passed in by **if**.

3.4.2 Recursion vs. iteration

In PCF, we introduce recursive definitions via the *fixpoint* special form that “computes” $F^\omega(\Omega) : T$ for $F : T \rightarrow T$ and so obviously violates affinity (in all but trivial cases).

$$\frac{\Gamma \vdash F : T \rightarrow T}{\Gamma \vdash (\mathbf{fixpt} F) : T}$$

Essentially, if F can access one of its free variables and then make a recursive call, we run the risk of reaccessing that same free variable in the new copy of F —whence $(\mathbf{fixpt} F)$ would violate affinity. But, if F is affine and moreover *never* accesses any of its free variables, $(\mathbf{fixpt} F)$ remains affine. We could thus add **fixpt** to affine PCF provided we only ever apply it to *closed* F .

Starting from the observation that accessing a free variable poses no problem, as long as we don’t subsequently make a recursive call, we could instead add a construct for *iteration*. This seems better adapted to the nature of affine PCF: an iterative process consists of some term M_i to build the initial state and another term M_t that describes the state transformation (together with the termination conditions). We could thus allow M_i to access as many free variables as it likes but require that M_t be closed. Such an iteration would be guaranteed to preserve affinity—but would, of course, require a richer type system, with products or lists, to be of any practical use.

4 Conclusions and further directions

We have shown that, in the setting of arena games, we can isolate the class of affine strategies. A number of questions now arise. Firstly, can we define the ! of linear logic as a comonad on this category to recover the original CCC of innocent strategies? It seems likely that the ! would have to operate (at the level of strategies) on *well-opened* plays, *not* on threads, to avoid the ambiguities otherwise inherent in plays as below.

$$\begin{array}{c} \mathbf{q}_1 \\ \backslash \\ \mathbf{t} \end{array} \quad \begin{array}{c} \mathbf{q}_2 \\ \backslash \\ \mathbf{t} \end{array}$$

Assuming that \mathbf{q}_1 and \mathbf{q}_2 are occurrences of different moves, should we think of this play as being one copy of a well-opened play (which happens to contain two initial occurrences) or as an interleaving of two copies of two well-opened plays (each of which actually happens to be a thread)?

A second interesting question concerns *nondeterministic* affine strategies. It has been known for some time that [with the “obvious” definition] nondeterministic innocent strategies fail to compose correctly, the failure of the *unique witness* property being the main culprit. It turns out that affine strategies reinstate this property, even in the absence of determinism: essentially, once two interactions between nondeterministic affine σ and τ have separated (in the “hidden” part of the interaction) they can never “come back together” and hence must witness different plays.

The resulting category of nondeterministic affine strategies still seems to be an SMCC and contains an erratic choice operator $\nabla : (A \times A) \rightarrow A$ which acts as a one-sided inverse to Δ (*i.e.* $\Delta_A ; \nabla_A = \text{id}_A$ in the ambient category \mathbf{G}). This category presumably models some kind of “affine erratic PCF”. However, it remains to be seen whether or not this sheds any light on the longstanding question of how to correctly define *innocent nondeterminism*, the problem essentially being how to make Δ and ∇ peacefully coexist in an innocent setting.

References

- [1] S. Abramsky, K. Honda, and G. McCusker. A fully abstract game semantics for general references. In *Proceedings, Thirteenth Annual IEEE Symposium on Logic in Computer Science*, 1998.
- [2] V. Danos and R. Harmer. The anatomy of innocence. In *Proceedings, Tenth Annual Conference of the European Association for Computer Science Logic*. Springer Verlag, 2001.

- [3] P. J. Freyd, P. W. O’Hearn, A. J. Power, M. Takeyama, R. Street, and R. D. Tennent. Bireflectivity. *Theoret. Comput. Sci.*, 228(1-2):49–76, 1999. Mathematical foundations of programming semantics (New Orleans, LA, 1995).
- [4] R. Harmer. *Games and full abstraction for nondeterministic languages*. PhD thesis, University of London, 1999.
- [5] J. M. E. Hyland and C.-H. L. Ong. On full abstraction for PCF: I, II and III. *Information and Computation*, 163:285–408, 2000.
- [6] P.-A. Melliès. Categorical models of linear logic revisited. *Theoretical Computer Science*, To appear.
- [7] H. Nickau. Hereditarily sequential functionals. In *Proceedings, Logical Foundations of Computer Science*, volume 813 of *Lecture Notes in Computer Science*. Springer-Verlag, 1994.