Events, Causality and Symmetry

Glynn Winskel
University of Cambridge Computer Laboratory, England
Glynn.Winskel@cl.cam.ac.uk

Abstract

The article discusses causal models, such as Petri nets and event structures, how they have been rediscovered in a wide variety of recent applications, and why they are fundamental to computer science. A discussion of their present limitations leads to their extension with symmetry. The consequences, actual and potential, are discussed.

Keywords: Petri nets, Event structures, Unfolding, Symmetry, Spans, Profunctors, Domain theory, Concurrency

1. INTRODUCTION

We are witnessing a rebirth of interest in causal models. Causal models are alternatively described in a variety of ways, as: causal-dependence models, independence models, non-interleaving models, true-concurrency models, and partial-order models. They include Petri nets, event structures, Mazurkiewicz trace languages, transition systems with independence, multiset rewriting, and many more. Fortunately this diversity is to a large extent only apparent. The models share the central feature that they represent processes in terms of the events they can perform, and that they make explicit the causal dependency and conflicts between events. Although this might be done through different mechanisms, *e.g.* the effect events have on local states (as is the case for Petri nets) or more abstractly through relations which directly express the causal dependency and the conflict/consistency of events (in event structures), the different models can be formally related—see Section 2.2.

Causal models have arisen, and have sometimes been rediscovered as *the* natural model, in many diverse and often unexpected areas of application:

Security protocols: for example, strand spaces are a form of event structure which support reasoning about secrecy and authentication through causal relations [30, 15];

Systems biology: biologists rediscovered Petri nets in the analysis of chemical pathways (with conditions standing for molecular species and events for reactions). Ideas from Petri nets and event structures are exploited in recent descriptive and analysis tools in tracking chemical pathways [58, 18];

Physics: as causal sets in theories of quantum gravity [64];

Hardware: in the design and analysis of asynchronous circuits [81];

Types and proof: as representations of propositions as types, and of proofs [26, 3, 17];

Nondeterministic dataflow: where numerous researchers have used or rediscovered causal models in providing a compositional semantics—see [55] and its references;

Network diagnostics: in the monitoring and fault diagnosis of communication networks [6];

Logic of programs: in concurrent separation logic where some artificiality in Brookes' pioneering soundness proof are obviated through a Petri-net model [31];

Partial order model checking: following the seminal work of McMillan [46] the unfolding of nets is exploited in the automated analysis of systems [20];

Distributed computation: event structures appear both classically [40] and recently in the Bayesian analysis of trust [51].

To illustrate the close relationship between Petri nets (based on the changes events incur on local states) and the 'partial-order models' of occurrence nets and event structures (possessing a global partial order of causal dependency on events), we consider how a Petri net can be unfolded first to a net of occurrences and from there to an event structure [52]. The unfolding construction

A Petri net and its unfolding

is analogous to the well-known method of unfolding a transition system to a tree, and is central to several analysis tools in the applications above. In the figure, the net on the left has loops. It is an example of a (1-safe) Petri net. Its conditions drawn as circles stand for local states. An event, a rectangle, when it occurs ends the holding of its preconditions (those conditions with arcs into the event) and begins the holding of its postconditions (those conditions with arcs from the event). Initially the two conditions at the bottom are imagined to hold, shown by their being marked. Initially any of the three events with marked preconditions can occur, ending the holding of its respective precondition and beginning the holding of its postcondition. Though of them, the two events on the left are in *conflict*, in the sense that only one of them can occur—they compete to end the holding of their common precondition. Either of those two events can occur *concurrently* with the third event to the right, in the sense that the third event shares no pre- or postconditions with them and so can occur independently. Once one of the two conflicting events and the event to the right have occurred all the preconditions of the top event will hold and it can occur, restoring the marking of conditions to its original state.

The net to the right is its *occurrence-net unfolding*. It consists of all the occurrences of conditions and events of the original net, and is infinite because of the original repetitive behaviour. The occurrences keep track of what enabled them. Notice how the shared postcondition of the conflicting events on left splits into two occurrences according to which of the two conflicting events gave rise to it. Similarly the top event splits into occurrences according to the nature of the occurrence of its left precondition.

The conditions in the occurrence net play two roles. They provide links of causal dependency between event occurrences. They also show when event occurrences are in conflict through sharing a common precondition. The simplest form of event structure arises by abstracting away the conditions in the occurrence net and capturing their two roles in relations of causal dependency and conflict on event occurrences.

Despite their often very successful, specialized applications, causal models lack a *comprehensive* theory which would support:

- their systematic use in giving *structured operational semantics* to a broad range of programming and process languages; while many examples exist of Petri-net semantics of processes and languages there are presently no generally-accepted standard techniques for describing operational semantics using Petri nets on similar lines to Plotkin's 'Structural Operational Semantics' [57].
- an expressive '*domain theory*' with rich higher-order type constructions needed by mathematical semantics. It should for example cover the standard event-structure semantics of CCS [70], extend to higher-order CCS, and support the formalization and analysis of distributed algorithms. Such a domain theory would go beyond traditional domain theory, in which types are represented as partial orders of information, in that causal models (perhaps

enriched, *e.g.* with probability) would feature as denotations. (The very concreteness of causal models appears to belie this possibility.)

This paper argues for a research programme towards such a comprehensive theory, and why a comprehensive theory, including a ‘domain theory,’ of causal models is within reach. A remedy to the overly-concrete nature of causal models is presented: a formal treatment of symmetry in causal models reveals connections between causal models *with symmetry* and the rich and expressive world of higher-dimensional algebra, in which theories of types, homotopy, geometry and combinatorics also find a mathematical home [23]. In fact much of the work reported here started in the search for an operational reading of specific semantics using higher-dimensional algebra [12, 54, 55] and the realization that denotations (as presheaves and profunctors) could sometimes be represented by event structures [53, 60].

As should be clear from the existing application areas the work on causal models would aim to contribute to several of the BCS ‘Grand challenges’: GC1 In vivo—in silico; GC2/4 Ubiquitous computing; GC6 Dependable system evolution; GC7 Journeys in nonclassical computation.

2. HISTORY

As we experience the ever-broader uses of computers so must we adapt our understanding of what a computational process is.

In the earliest days of computer science it became accepted that a computation was essentially an (effective) partial function

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

between the natural numbers. This view underpins the Church-Turing thesis on the universality of computability.

As computer science matured it demanded increasingly sophisticated mathematical representations of processes. The pioneering work of Strachey and Scott in the denotational semantics of programs assumed a view of a process still as a function

$$f : D \rightarrow D' ,$$

but now acting in a continuous fashion between datatypes represented as special topological spaces, ‘domains’ D and D' ; reflecting the fact that computers can act on complicated, conceptually-infinite objects, but only by virtue of their finite approximations.

In the 1960’s, around the time that Strachey started the programme of denotational semantics [65], Petri advocated his radical view of a process, expressed in terms of its events and their effect on local states [44]—a model which addressed directly the potentially distributed nature of computation, but which, in common with many other current models, ignored the distinction between data and process implicit in regarding a process as a function. I will argue that today a serviceable notion of process requires a marriage of Petri’s view of a process and the vision of Scott and Strachey.

2.1. Classical domain theory

It is helpful to recall the basics of traditional domain theory, the mathematical foundations of denotational semantics. A *domain* is a partial order $(D, \sqsubseteq)$ with, at the very least, the completeness property that any infinite chain

$$d_0 \sqsubseteq d_1 \sqsubseteq \cdots \sqsubseteq d_n \sqsubseteq \cdots$$

has a least upper bound $\bigsqcup_n d_n$. The order $\sqsubseteq$ can be mysterious to a beginner and is quite abstract. It stands for an order of increasing computational information, information which can be presented as a limit (the least upper bound) of a chain of approximations. Accordingly, a function between

domains $f : D \rightarrow E$ should be *continuous* in that f should preserve limits, *i.e.* it should preserve the information order and least upper bounds of chains.

If a domain D has a least element $\perp$ and a function $f : D \rightarrow D$ is continuous, then f has a least fixed point given by $\bigsqcup_n f^n(\perp)$. This is a central tool in giving meaning to recursive programs. By pushing this technique for solving recursive definitions up to the level of domains (treating the category of domains with continuous functions as analogous to a domain itself), Scott achieved the breakthrough in the late sixties of producing a nontrivial solution to $D \cong [D \rightarrow D]$ (a recursively defined domain), so providing a model of the λ -calculus, and, by the same techniques, the semantics of recursive types [62].

Representations of domains

What is the information order? Can any sense be made of its ‘units’ of information? There are essentially two answers in the literature, the ‘*topological*,’ the most well-known from Scott’s work, and the ‘*temporal*,’ arising from the work of Berry [7]:

Topological: the basic units of information are *propositions* describing finite properties; more information corresponds to more propositions being true. Functions are ordered pointwise. Domains are represented by logical theories in the form of ‘information systems’ or ‘logic of domains.’

Temporal: the basic units of information are *events*; more information corresponds to more events having occurred over time. Functions are restricted to ‘stable’ functions and ordered by the intensional ‘stable order,’ in which common output has to be produced for the same minimal input. Berry’s specialized domains ‘dl-domains’ are represented by event structures.

In truth, Berry developed ‘stable domain theory’ by a careful study of how to obtain a suitable category of domains with stable rather than all continuous functions. He arrived at the axioms for his ‘dl-domains’ because he wanted function spaces (so a cartesian-closed category). The realization that dl-domains were precisely those domains which could be represented by event structures, a fact I shall explain now, came a little later [72, 75].

Event structures

An *event structure* comprises $(E, \leq, \text{Con})$, consisting of a set E , of *events* (event occurrences), partially ordered by $\leq$, the *causal dependency relation* which satisfies $\{e' \mid e' \leq e\}$ is finite for all $e \in E$, and a family Con of finite subsets of E , the *consistency relation*, which satisfy

$$\begin{aligned} \{e\} &\in \text{Con} \text{ for all } e \in E, \\ Y \subseteq X \in \text{Con} &\implies Y \in \text{Con}, \text{ and} \\ X \in \text{Con} \ \& \ e \leq e' \in X &\implies X \cup \{e\} \in \text{Con}. \end{aligned}$$

The relation $e' \leq e$ expresses that the occurrence of the event e depends on the previous occurrence of event e' . The consistency relation picks out those events which can occur together.

Consider a record of the events that have occurred. If an event e has occurred then so must all events e' on which e depends have occurred previously, and any finite set of events that have occurred should be consistent. Accordingly, the *configurations*, $\mathcal{C}(E)$, of an event structure E consist of those subsets $x \subseteq E$ which are

$$\begin{aligned} \text{Down-closed: } &\forall e, e'. \ e' \leq e \in x \implies e' \in x \quad \text{and} \\ \text{Consistent: } &\forall X \subseteq_{\text{fin}} x. \ X \in \text{Con}. \end{aligned}$$

Configurations stand for histories, given in terms of the events that have occurred; the events inherit an order from the event structure. In particular, for an event e the set $[e] =_{\text{def}} \{e' \in E \mid e' \leq e\}$ is a configuration including the whole causal history of the event e . The inclusion relation $x \subseteq x'$ means that x is a sub-history of x' . Ordered by inclusion the configurations form

a domain $(\mathcal{C}(E), \subseteq)$ —in fact, when E is countable, a dl-domain as discovered by Berry, and all such are so obtained. In this way event structures can represent a rich variety of types, even for polymorphism [13, 75].

Anomalies in domain theory

Nondeterminism: For traditional ('topological') domain theory the problem of adjoining nondeterminism was solved by Plotkin through the introduction of powerdomains [56]. Powerdomains can be seen as concerned with information about the possible and eventual properties of a nondeterministic process [73]. But for stable domain theory the information order of all but the simplest powerdomain fail to be temporal. Section 2.3 will provide an alternative way to adjoin nondeterminism.

Concurrency/interaction: The intricacy of models for distributed computation such as Petri nets and event structures (modelling both processes and types), means that they don't fit comfortably within a partial order of information. Rather their intricacy suggests that they belong more rightfully to an extended notion of domain as a category—see Section 2.2. Only rarely do the wanted equivalences on processes arise from traditional domain theory.

Probability and nondeterminism: There are powerdomains both for nondeterminism and probability. Sometimes one needs both. Combining probability and nondeterminism is problematic because the two forms of powerdomain together do not satisfy a distributive law (their combination forces extra laws to be imposed [48]). However, if one works with the *indexed probabilistic powerdomain* where the probability distribution is carried by the ways in which values are computed, one recovers a distributive law [68, 67].

Nondeterministic dataflow: While, as Kahn was early to show [39], deterministic dataflow is a shining application of simple domain theory, nondeterministic dataflow is beyond its scope. The compositional semantics of nondeterministic domain theory needs a form of generalized relation which specify the ways input-output pairs are realized—see Section 2.3.

In my view the problems point towards a more intensional 'domain theory' which expresses the set of ways of computing. Early suggestions along these lines were made by Lehman with a definition of a categorical powerdomain [41] and by Girard in his 'quantitative' domain theory [25]. Both fit within the broader view of domains as presheaf categories [12]. Causal models have reappeared in providing an operational reading of the ways computations are realized, for example as the finite configurations of an event structure [53].

2.2. An abundance of models

Partly in reaction to the difficulties of traditional denotational semantics and domain theory, today we find a range of ways to model a process in computer science, for example as: a transition system; an equivalence class of transition systems w.r.t. some equivalence such as bisimulation; a coalgebra; a resumption in a powerdomain; an IO-automaton; a Petri net; a game; a (generalized) relation; . . .

Fortunately many models can be formally related by adjunctions whose adjoints give translations of one model into another. The adjunctions make explicit how to translate from one kind of model (say Petri nets) to another (say occurrence nets, or event structures); this relies on regarding a kind of model (say Petri nets) as a category (a category of Petri nets) with suitable 'simulation' maps (we define the maps on event structures below). For example the unfolding of a Petri net illustrated in the Introduction is the right adjoint to the inclusion functor from the category of occurrence nets to the category of (1-safe) Petri nets. There is an intuitively obvious map $f : \mathcal{U}(N) \rightarrow N$ from the occurrence net $\mathcal{U}(N)$ back to the original net N ; it takes occurrences of conditions and events back to those conditions and events in the original net of which they are occurrences. A way to express the adjunction is through the following *universal characterization* of the unfolding. Given any map $g : O \rightarrow N$ from an occurrence net O to the original net, there is a unique map $h : O \rightarrow \mathcal{U}(N)$

such that $f \circ h = g$:

$$\begin{array}{ccc} \mathcal{U}(N) & \xrightarrow{f} & N \\ \uparrow h & \nearrow g & \\ O & & \end{array}$$

[For a definition of the maps of Petri nets, see [74, 70].]

As a right adjoint, unfolding automatically preserves limits, for example products and pullbacks, in the category, and this can be useful in relating parallel compositions in one model to those in another [21]. There is a further adjunction between event structures and occurrence nets; its right adjoint ‘strips’ away the conditions of an occurrence net to reveal its underlying event structure, while its left adjoint ‘saturates’ an event structure with conditions to make an occurrence net. Adjunctions compose so we obtain an adjunction between event structures and safe nets—its right adjoint acts as the operation from Petri nets to event structures sketched in the Introduction.

The use of categories exposes a uniformity across different models. Semantics of synchronising processes [34, 47], whether they be in transition systems, Petri nets, event structures or many other models, are given in precisely the same way in terms of the categorical constructions used. Presented as categories, models support a general, diagrammatic definition of an important equivalence, *strong bisimulation* and its extension beyond transition systems, via open maps—see Section 3.2.

The categories depend on a choice of simulation map which we illustrate for event structures. The maps arise in relating compound constructions to their components, for example, as projections from a parallel composition of event structures to one of its components. Earlier we saw event structures as representations of domains, so types. Now we are seeing them in the role of processes (in the sense of the synchronising processes of Hoare and Milner in CSP and CCS). Let E and E' be event structures. A *partial map* of event structures $f : E \rightarrow E'$ is a partial function on events $f : E \rightarrow E'$ such that for all $x \in \mathcal{C}(E)$

$$\begin{aligned} &fx \in \mathcal{C}(E') \text{ and} \\ &\text{if } e_1, e_2 \in x \text{ and } f(e_1) = f(e_2), \text{ both being defined, then } e_1 = e_2. \end{aligned}$$

The idea is that the occurrence of an event e in E induces the *coincident* occurrence of the event $f(e)$ in E' , whenever it is defined. Because events are thought of as essentially without duration, two distinct events which occur in a history of the input cannot be coincident with a common event in the image—the reason for the ‘local injectivity’ condition.

Despite the abundance of categories of pre-existing models, they form a rather patchy landscape and are, for example, insufficient to represent *higher-order* processes (which might take a process itself as input and deliver another process as output). Presheaf categories fill out the landscape of models to provide a versatile range of models for processes [12]. The idea is to build models for processes directly out of computation paths, regarding a nondeterministic process as a presheaf on a category of paths; essentially, a presheaf is a glueing together of computation paths. Presheaf categories are as versatile as the notion of computation path. With suitable choices of paths, presheaf categories subsume existing models such as event structures, while supporting a range of type constructions, also for higher-order processes and name generation [76, 12, 53, 11]. Presheaf categories and the relations between them, expressed as profunctors, connect with the rich world of higher-dimensional algebra. In particular, the little-explored representation of processes as ‘bundles’ is crucial in the general treatment of weak bisimulation, and its extension to causal models [9, 22].

2.3. Relations rediscovered

From a historical perspective it is remarkable that so many of the models of processes listed in Section 2.2 are not associated with input and output types; while the structure of models has

become much more sophisticated the explicit connections with input and output data have often been lost.

A central exception is that of generalized relations in the form of *profunctors* (or *distributors*, or *bimodules*). Profunctors are relations between categories in which the category of sets takes over the role of truth values; instead of simply saying whether or not an object of input is related to an object of output, a profunctor provides a set of ways in which that particular input-output instance is realized. Technically, a profunctor F from a category A to a category B , written $F : A \multimap B$, is a functor $F : A \times B^{\text{op}} \rightarrow \text{Set}$, covariant in input and contravariant in output (see [77] for more intuition).

Profunctors have arisen independently in a range of areas: in logic and types, e.g. Girard's normal functors [25] and 'container types' [1]; combinatorics through Joyal's theory of species [37] and its extensions [23]; nondeterministic dataflow [55]; higher-order programming languages and processes [12]; categories of models for concurrent computation [12]; as a starting point for the theory of operads [16]; ...

The 'relations' arising from computation can often be represented, in a more computationally informative way, in terms of event structures, with event structures playing both the role of input and output *types*, as well as the *process* of computation between them. A compelling example comes from the early work of Brock and Ackerman who were the first to emphasize the difficulties in giving a compositional semantics to nondeterministic dataflow [8], though our example is based on simplifications in the later work of Rabinovich and Trakhtenbrot, and Russell.

Nondeterministic dataflow—the Brock-Ackerman anomaly

There are two simple nondeterministic processes A_1 and A_2 , which have the same input-output relation, and yet behave differently in the common feedback context $C[-]$, illustrated above. The context consists of a fork process F (a process that copies every input to two outputs), through which the output of the automata A_i is fed back to the input channel, as shown in the figure. Process A_1 has a choice between two behaviours: either it outputs a token and stops, or it outputs a token, waits for a token on input and then outputs another token. Process A_2 has a similar nondeterministic behaviour: Either it outputs a token and stops, or it waits for an input token, then outputs two tokens. For both automata, the input-output relation relates empty input to the eventual output of one token, and non-empty input to one or two output tokens. But $C[A_1]$ can output two tokens, whereas $C[A_2]$ can only output a single token. Notice that A_1 has two ways to realize the output of a single token from empty input, while A_2 only has one. It is this extra way, not caught in a simple input-output relation, that gives A_1 the richer behaviour in the feedback context.

Over the years there have been many solutions to giving a compositional semantics to nondeterministic dataflow (see [55] for fuller references). But they all hinge on some form of generalized relation, to distinguish the different ways in which output is produced from input. A compositional semantics can be given using *stable spans* of event structures, an extension of Berry's stable functions to include nondeterminism [60, 61]. A process of nondeterministic dataflow, with input type given by an event structure A and output by an event structure B , is captured by a pair of maps (a *span*)

where E is also an event structure. The map $out : E \rightarrow B$ is a *rigid* map, *i.e.* a total map of event structures as in Section 2.2 which preserves the relation of causal dependency, or equivalently, a total map with the property that for a configuration x of E if y is a subconfiguration of $out\ x$ then there is a (necessarily unique) subconfiguration x' of x such that $out\ x' = y$:

$$\begin{array}{ccc} x' & \subseteq & x \\ \downarrow & & \downarrow \\ y' & \subseteq & out\ x \end{array}$$

The map $dem : E \rightarrow A$, associated to input, is of a different character. It is a *demand* map, *i.e.* a function from $\mathcal{C}(E)$ to $\mathcal{C}(A)$ which preserves finite configurations and unions; $dem\ x$ is the minimum input for x to occur and is the union of the demands of its events. The occurrence of an event e in E demands minimum input $dem[e]$ and is observed as the output event $out(e)$. *Deterministic* stable spans, where consistent demands in A lead to consistent behaviour in E , correspond to Berry's stable functions.

The stable span determines a *profunctor* $\tilde{E}$ from the finite configurations p of A to the finite configurations q of B :

$$\tilde{E}(p, q) = \{x \in \mathcal{C}(E) \mid dem\ x \subseteq p \ \& \ out\ x = q\},$$

the set of ways the input-output pair (p, q) is realized.

Stable spans can be composed one after the other (essentially by a pullback construction, as rigid maps extend to special demand maps between configurations)—their composition coincides with the composition of their profunctors. They also have a nondeterministic sum, and compose in parallel, and most significantly allow a feedback operation [60]. In fact, stable spans were first discovered as a way to represent, and give operational meaning to, the profunctors that arose as denotations of terms in affine-HOPLA, an affine Higher Order Process LAnguage [53, 54]. The spans helped explain the tensor of affine-HOPLA as the parallel juxtaposition of event structures and a form of entanglement which appeared there as patterns of consistency and inconsistency on events. The use of stable spans in nondeterministic dataflow came later as a representation of the profunctors used in an earlier semantics [55, 60].

3. CAUSAL MODELS AND SYMMETRY

3.1. Anomalies in traditional causal models

The insufficiency of stable spans: Although one can easily encode CCS within affine-HOPLA this induces the usual interleaving semantics of CCS. One can prove that the spans denotable by terms of affine-HOPLA can never be those of the usual event-structure semantics of CCS [70]. More generally, because output maps of stable spans are rigid, they are too restrictive when considering parallel compositions via synchronizations. Such a parallel composition of event structures can augment extra causal dependency to that which is present in the event structures originally. So 'projections' from a parallel composition to a component are rarely rigid.

Varying maps: There are uses for several different forms of maps on event structures: rigid, total and partial maps, demand maps and a variant of demand maps due to Abbes [2]. Changing the category generally changes important categorical constructions. One would like to settle on some basic maps and then have a systematic way to vary the nature of maps within it.

Unfoldings of general Petri nets: In general nets conditions can hold with multiplicities. While their occurrence net unfoldings can be defined, there is no universal characterisation like that of Section 2.2. The symmetry intrinsic to nets with multiplicities spoils uniqueness.

Unfoldings of higher-dimensional automata: Higher-dimensional automata are essentially glueings together of cubes of concurrent actions [29]. Although they can be unfolded to event structures the identifications due to glueing spoil uniqueness in trying to get a universal characterization.

Weak bisimulation: Just as for labelled transition systems weak bisimulation between labelled event structures (in which we abstract from invisible actions, generally labelled τ) can be explained as strong bisimulation between the results of ‘hiding’ the invisible actions [22]. Whereas the ‘hiding’ operation on a transition system is again a transition system, the hiding operation on an event structures does not always yield an event structure [22].

Name generation: There are methods to represent the generation of new names in causal models, e.g. [15, 77]. But the methods are overly concrete in the sense that they ignore the implicit symmetry on names. Presently there are difficulties in extending work on an event-structure semantics of the pi-Calculus [14] to the whole language because of the absence of a key algebraic operation—a form of new-name abstraction on event structures.

The anomalies have a common solution: a formal treatment of symmetry in processes. Consider the first, that stable spans are insufficient. One wishes to free up the choice of maps for the input and output legs of the span. In answer to the second difficulty one would like to do this systematically.

A systematic way to modify maps is through Kleisli maps associated with monads [49]. Starting from an original category—a good choice would be the category of event structures with rigid maps—one could hope to obtain other kinds maps from an object E to an object B as original maps from E to $T(B)$, where T is an appropriate monad. This suggest that stable spans be generalized to general spans of event structures

with input the event structure A , output B and process E , w.r.t. suitable monads S and T to moderate the regimes of input and output. More should hold for the spans to compose [78, 79].

In the span shown above, one would hope in particular for a monad S such that demand maps from E to A are realized as Kleisli maps from E to $S(A)$. It becomes important that event structures are able to support a reasonable repertoire of monads. It is here we run into difficulties. Consider for example the useful operation of replication $!E$ forming the parallel composition of countably many copies of an event structure E . For this to be a monad we require a unit $\eta_E : E \rightarrow !E$ and multiplication $\mu_E : !!E \rightarrow !E$. It would seem reasonable that η_E takes E to the zeroth copy. Assuming an injection $[-, -] : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ encoding pairs of natural numbers as natural numbers, again it would seem reasonable that μ_E takes $(i, (j, e))$ to $([i, j], e)$. But then the laws for monads fail. It is easy to check that none of the monad diagrams commute:

For example if the first two diagrams commuted, $i = [0, i] = [i, 0]$, violating the injectivity of pairing. But in $!E$ one copy of E is similar to another. Up to this symmetry, allowing one copy to swap with another, the diagrams do commute. In answer to first and second anomalies at least, a formal treatment of symmetry is needed.

3.2. Symmetry

A treatment of symmetry on models makes use of a general method of open maps in defining bisimulation in a variety of models. The method shows one advantage of presenting models as categories—see [38] for more details.

Let $\mathcal{C}$ be a category with a distinguished subcategory $\mathcal{P}$ of path objects with path-extension maps. For example, the category $\mathcal{C}$ might be that of event structures and $\mathcal{P}$ the subcategory of finite

‘elementary’ event structures in which all subsets of events are consistent—so partial-order runs. Alternatively, the category $\mathcal{C}$ could be the category of transition systems and $\mathcal{P}$ the subcategory of special transition systems that have the form of single finite branches. (We might also allow infinite runs, important in treating ‘fairness.’) A map $f : X \rightarrow Y$ in $\mathcal{C}$ is a bisimulation map, traditionally called *open*, if, for any map $s : P \rightarrow Q$ in $\mathcal{P}$ and maps $p : P \rightarrow X$ and $q : Q \rightarrow Y$, if the diagram

$$\begin{array}{ccc} P & \xrightarrow{p} & X \\ s \downarrow & & \downarrow f \\ Q & \xrightarrow{q} & Y \end{array}$$

commutes then there is a map $h : Q \rightarrow Y$ such that the diagram

$$\begin{array}{ccc} P & \xrightarrow{p} & X \\ s \downarrow & \nearrow h & \downarrow f \\ Q & \xrightarrow{q} & Y \end{array}$$

commutes, i.e. $p = hs$ and $q = fh$. This path-lifting property says that any extension, viz. qs of a path fp in Y can be matched, via f , by an extension hs of the path p in X . It can be shown by straightforward diagrammatic arguments that open maps include all isomorphisms, are closed under composition (and therefore form a subcategory). It is an instructive exercise to show that open maps are preserved under pullbacks [38].

We define a bisimulation between two objects X and Y in $\mathcal{C}$ to be a generalized relation in the form of a span of open maps in $\mathcal{C}$:

$$\begin{array}{ccc} & R & \\ \swarrow & & \searrow \\ X & & Y \end{array}$$

When $\mathcal{C}$ has pullbacks, and categories such as those of event structures and transition systems do, bisimulation induces an equivalence relation on objects of $\mathcal{C}$ (using the fact that a pullback of an open map is open). In the case of transition systems the equivalence obtained is the usual strong bisimulation of Milner and Park [47], while for event structures it is also independently known, and called hereditary history-preserving bisimulation (unfortunately!) [5].

We now show how to extend a category of models to a category of models with symmetry. We won’t describe this in fullest generality, but for the case in which $\mathcal{C}$ is a category with pullbacks with distinguished subcategory $\mathcal{P}$ of path objects.

We define a new category, $\mathcal{SC}$, of objects with symmetry and symmetry-preserving maps, as follows:

The *objects* of $\mathcal{SC}$ are triples $(C; l, r : S \rightarrow C)$ comprising an object C of $\mathcal{C}$ together with *open* maps $l : S \rightarrow C$ and $r : S \rightarrow C$, from a common object S in $\mathcal{C}$, which form a pseudo equivalence (see the Appendices for the definition of *open* map and *pseudo equivalence*).

An object $(C; l, r : S \rightarrow C)$ of $\mathcal{SC}$ is an object of $\mathcal{C}$ with a span of open maps

$$\begin{array}{ccc} & S & \\ l \swarrow & & \searrow r \\ C & & C \end{array}$$

to be thought of as expressing when paths in C are related by symmetry. We call objects of $\mathcal{SC}$ *objects with symmetry*. When l, r form an equivalence relation we shall call $(C; l, r : S \rightarrow C)$ a *symmetry equivalence*. In writing them we shall adopt the convention that for instance $(A; S_A)$ describes the object with symmetry $(A; l_A, r_A : S_A \rightarrow A)$.

In $\mathcal{SC}$, a *map* $f : (A; S_A) \rightarrow (B; S_B)$ is a map $f : A \rightarrow B$ in $\mathcal{C}$ which preserves symmetry in the sense that

$$\begin{array}{ccccc} A & \xleftarrow{l_A} & S_A & \xrightarrow{r_A} & A \\ f \downarrow & & h \downarrow & & \downarrow f \\ B & \xleftarrow{l_B} & S_B & \xrightarrow{r_B} & B \end{array}$$

commutes for some $h : S_A \rightarrow S_B$ in $\mathcal{C}$. Maps compose as maps in $\mathcal{C}$ and share the same identity maps.

For maps $f, g : (A; S_A) \rightarrow (B; S_B)$, define $f \sim g$ iff there is a map $h : A \rightarrow S_B$ in $\mathcal{C}$ such that

$$\begin{array}{ccc} & A & \\ f \swarrow & \downarrow h & \searrow g \\ B & \xleftarrow{l_B} S_B \xrightarrow{r_B} & B \end{array}$$

commutes for some $h : S_A \rightarrow S_B$ in $\mathcal{C}$.

The relation $\sim$ is an equivalence relation on maps $\mathcal{SC}(\mathcal{A}, \mathcal{B})$ between objects with symmetry A and B . The relation $\sim$ respects composition in the sense that if $f \sim g$ then $h \circ f \circ k \sim h \circ g \circ k$, for composable maps h and k . In other words, the category $\mathcal{SC}$ is enriched in the category of equivalence relations.

Maps f and g in $\mathcal{SC}$ for which $f \sim g$ are thought of as *the same up to symmetry*.

Equivalence on maps induces equivalence on objects: Let A and B be objects with symmetry. An *equivalence* from A to B is a pair of maps $f : A \rightarrow B$ and $g : B \rightarrow A$ in $\mathcal{SC}$ such that $f \circ g \sim \text{id}_B$ and $g \circ f \sim \text{id}_A$; then we say A and B are equivalent and write $A \simeq B$.

Under minor conditions, functors and adjunctions between categories with pullbacks and subcategories of paths lift to their extensions with symmetry.

Examples (1) If the subcategory $\mathcal{P}$ of $\mathcal{C}$ is a groupoid (i.e. all its maps are isomorphisms), then all maps of $\mathcal{C}$ are open and the construction of $\mathcal{SC}$ from $\mathcal{C}$ coincides with its *exact completion* [10].

(2) For the categories of event structures (with partial, total or rigid maps), a symmetry equivalence on an event structure E corresponds to an *isomorphism family* S of bijections $\theta : x \cong_S y$ between finite configurations such that

(i) $\text{id}_x : x \cong_S x$ for finite $x \in \mathcal{C}(E)$; if $\theta : x \cong_S y$, then $\theta^{-1} : y \cong_S x$;
if $\theta : x \cong_S y$ and $\varphi : y \cong_S z$, then $\varphi \circ \theta : x \cong_S z$.

(ii) if $\theta : x \cong_S y$ and $x' \subseteq x$ with $x' \in \mathcal{C}(E)$, then
 $\theta' : x' \cong_S y'$ and $y' \subseteq y$ for some (unique) $y' \in \mathcal{C}(E)$.

(iii) if $\theta : x \cong_S y$ and $x \subseteq x'$ for finite $x' \in \mathcal{C}(E)$, then
 $\theta' : x' \cong_S y'$ and $y \subseteq y'$ for some $y' \in \mathcal{C}(E)$.

Note that an event structure with symmetry now represents a *category* of finite configurations where maps are got by composing isomorphisms from its isomorphism family with inclusions; by (ii) all maps can be normalized to an isomorphism from the isomorphism family followed by an inclusion. There are characterizations in terms of isomorphism families of preservation of symmetry by maps of event structures and the relation $\sim$ on symmetry-preserving maps [78, 79].

(3) For transition systems (with the standard maps preserving initial states and transitions [70]) a symmetry equivalence on a transition system corresponds to an equivalence relation on the states of the transition system which is also a strong bisimulation.

3.3. Consequences of symmetry

Recent and ongoing work

A major consequence is that many wished-for monads on event structures, such as replication, do indeed become monads up to symmetry (technically forms of pseudo monad) once event structures are extended with symmetry. Starting with the category of event structures with rigid maps, its extension with symmetry has monads with which to realize demand maps (including those of [2]), total non-rigid maps, partial maps and replication [78, 80]. The monads adjust the notion of event, including atomicity, so that events can now have duration. For example, the monad for demand maps creates events in the form of histories, describing the way that input is explored.

This opens the way to general spans providing semantics to potentially rich process languages and event types, supporting case analysis on events. This work is incomplete, but see [79] for an example of how such a higher-order language can induce the usual event-structure semantics for CCS, as well as an event-structure semantics for a higher-order variant of CCS.

By adjoining symmetry to Petri nets we can obtain a universal characterisation of unfoldings of *general* Petri nets, like that of Section 2.2, but where instead of uniqueness we achieve uniqueness up to symmetry [32]. (An explicit notion of symmetry for Petri nets was also used in [59].)

By extending event structures with symmetry we are able to represent a broader category of presheaves. Operations that previously only worked on presheaves now work on event structures with symmetry. We can now obtain a universal characterisation of an unfolding of higher-dimensional automata as an event structure with symmetry [79]. Now, in principle, the hiding operation on event structures with symmetry yields an event structure with symmetry, bringing a central operation of weak bisimulation back into causal models (this relies on a stronger result than that reported in [79]). Ongoing work on a domain theory [66] within the theory of nominal sets [43], was designed with the idea of extending it to presheaf models. Viewing event structures as presheaves (now in nominal sets) would yield constructs such as new-name abstraction on event structures, the missing key to an algebraic treatment of name generation in causal models.

In summary, the introduction of symmetry brings a new expressive power to causal models: previously unthought-of semantics to higher-order types and languages; mechanisms for event abstraction, allowing the switch from atomic to compound events; extensions to the usual Petri-net unfolding and its preservation results; and possibilities of new enrichments with further structure through the presentation of generalized relations as spans of causal models.

RESEARCH AREAS

So far the article has mainly shown how causal models and the introduction of symmetry arise from mathematical concerns, one thing leading to another, so to speak. But of course the importance of causal models to computer science derives from their wide range of application, and the numerous research areas in which they have potential significance. In the past it has been not so much the mathematical richness of causal models, but their almost physical nature, making them close to what is modelled or an implementation, that has most influenced their use. But it is their under-appreciated and under-explored mathematical richness that is likely to drive their new uses in the research areas below.

Domain theory for causal models: The explanation and development of causal models here, in particular the introduction of symmetry, have been motivated by considering how causal models might support a form of ‘domain theory’ and denotational semantics, though of a revolutionary form in which causal models play the role of both types and maps (in the form of spans). A lot remains to be done—see [79].

Distributed algorithms: One motivation in boosting the power of causal models has been to make distributed algorithms amenable to semantic description and reasoning techniques. The

potential richness of event types and higher-order languages should, when developed more fully, be helpful in the design and analysis of distributed algorithms. For a number of years now the techniques of distributed algorithms have been relevant in chip design, recognized in the use of Petri nets there [81], and becomes increasingly so as chip design is forced to count the cost of communication [50]. Significant advances have been made in the more limited regime of security protocols, but there are gaps in the treatment and exploitation of symmetry, and in relations with true cryptography, which calls for probability. Two other specialized areas which rely on causal models, and probabilistic event structures [69], are the distributed diagnosis of communication networks [6], and the very recent Bayesian analysis of event-based trust [51].

Probability and nondeterminism: The use of spans suggests largely unexplored ways to enrich computation with probability, in addition to nondeterminism, essentially by taking the vertex in a span of event structures to be a probabilistic event structure.

Systems biology: Causal models, with stochastic information, are already used in bio-computation [33] and there is reason to expect that higher-order bio-processes (accessible to the theory here) will become an important abstraction tool. Issues of compound versus atomic events are important in systems biology. In the presence of symmetry there are ways to realize forms of compound event, for example the 'persistent' events of [77], out of atomic events. Interestingly biologists are not interested in quite the same unfolding of Petri nets to event structures as computer scientists. Biologists are more interested in the 'macroscopic' events of a biological process rather than the 'microscopic' events which are picked out by the traditional computer-science unfolding. For example, in the traditional unfolding of a Petri net the participation of the *same* molecule at several stages in a reaction would lead to a chain of causal dependencies. However to a biologist that one particular molecule rather than another has participated in a reaction is generally neither here nor there, and such incidental dependencies would be dropped [18].

Weak memory models: Weak memory models [36] are challenging as they do not respect a property fundamental of most causal models, that events can be globally serialized. A memory transaction can be a compound event, comprising several different events, and which constituent event is in view depends on the processor. The relation between weak memory models and their implementation involves event abstractions.

Unfoldings and tools: Unfoldings of Petri nets have provided strikingly successful methods for the analysis of distributed systems [20]. The characterization of the unfolding as a right adjoint has been exploited in network diagnosis [21]. The techniques here based on symmetry extend unfoldings and their characterization, and begin to push Petri nets into the new territory of higher-order processes (although we have concentrated on event structures, similar ideas are working for nets). The compact, often finite, representations afforded by Petri nets accede to the known decision procedures for Petri nets, and regular languages and so potentially tools for the analysis of higher-order processes (in the manner of algorithmic game theory[24]). Symmetry is already exploited in model checking [19].

Structured Petri-net semantics and logic: There is the practical issue of structural operational semantics for Petri nets. In several separate Petri-net semantics, that of a language for security protocols [15], a language for biochemical systems [58], and a language with parallel commands and semaphores [31], the same technique has been used, and seems to be much more widely applicable. In defining a Petri net (or its abbreviation as a form of coloured Petri net), first its basic conditions for data, names, resources and control are defined followed by definition of its events in terms of their pre- and postconditions. The definitions proceed in a syntax-directed way, much as in Plotkin's SOS, but where in SOS it is generally transitions which are given inductively, here it is events. One feature currently missing from such semantics is that of symmetry, for example between the different but similar conditions standing for names of resources.

Quantum systems and event structures: There are intriguing parallels between quantum processes and processes described as event structures, where a degenerate form of entanglement appears in patterns of consistency and conflict amongst events, labels on copies

of identical processes behave like amplitudes, and ‘probabilistic tests’ [69] have much in common with the ‘consistent families’ of the ‘consistent-histories’ approach to quantum theory. Curiously the structure of consistency on events, present in event structures, is missing in causal sets. (A specific challenge, which may involve ideas from event structures, is to provide a denotational semantics for the full Selinger-Valiron lambda calculus for quantum computation [63].)

An event structure arises quite naturally in describing the histories of a quantum system (with states in Hilbert space H). This (untried) construction is inspired by the ‘consistent-histories’ approach to quantum theory [28]. A quantum history of the quantum system might be expressed as a finite sequence of quantum properties or observations, taken to be projectors on H . The quantum histories possess the structure of a Mazurkiewicz trace language by interpreting symbols of the alphabet as the quantum properties, and Mazurkiewicz-independence as commutability of distinct projectors. (It seems sensible to restrict to sequences whose composition is nonzero.) By a standard construction we can obtain an event structure from the Mazurkiewicz trace language [70]. It will be an event structure where the events are labelled by the quantum properties, so events correspond to occurrences of these observations; moreover, each finite configuration will be associated with an operator, that got by composing the sequence of projectors from which the configuration arose. With the help of the association of operators with finite configurations we can specify those sets of configurations to which it would be reasonable to ascribe a probability distribution. Such a set should be a maximal subset of finite configurations which pairwise have orthogonal operators (such sets of configurations are essentially ‘consistent histories’ of the ‘consistent-histories’ approach to quantum theory, though of course the use of ‘consistent’ here differs from that in event structures).

Continuous processes: A dream in the early development of event structures (see [71]) was to extend them to continuous processes so opening up their application to physical processes—Nachbin’s text was inspirational [42], with ideas of approximating continuous by discrete processes, but the enterprise dried up because of a lack of motivating examples. A revitalization has come through recent work showing how to reconstruct spacetime topology from causal structure using purely order-theoretic ideas [45].

Mathematics and logic: The research cannot be divorced from mathematical questions. Event structures represent types, underlie game semantics [4, 35] and begin to appear as denotations of proofs [3, 17]; the need for symmetry also appears here and in the earlier but related geometry of interaction of Girard [27]. It is not clear if the bipartite nature of the games (‘opponent’ vs. ‘player’) and the often intricate structure associated with it can be accommodated within spans of causal models. Perhaps it can. There are, for example, ways to express the composition of sequential algorithms as relational composition and this is taking us close to the composition of spans. The connection with higher-dimensional algebra has already been fruitful [12, 23] and is likely to lead to the refinement and development of causal models, as we seek to give computational interpretation to mathematical constructions. We cannot expect to separate causal models from the more purely mathematical structures of higher-dimensional algebra. In particular, we need to understand better the relationship between the two kinds of generalized relations, spans of event structures and (certain kinds of) profunctors.

Acknowledgements

I have benefited from discussions with Marcelo Fiore, Jon Hayman, Martin Hyland, Jean Krivine, Prakash Panangaden, Gordon Plotkin, Lucy Saunders-Evans, Pawel Sobocinski, Peter Sewell and Sam Staton, and the advice of the anonymous referees.

REFERENCES

- [1] Abbott, M., Altenkirch, T., Ghani, N., Categories of containers. FSSCS’03, LNCS 2620, 2003.
- [2] Abbes, S., A Cartesian closed category of event structures with quotients. Discrete Mathematics and Theoretical Computer Science (DMTCS) 8(1):249-272, 2006.
- [3] Abramsky, S., Event Domains, Stable Functions and Proof Nets. ENTCS 172, 33–67, 2007.

- [4] Abramsky, S., Jagadeesan, R., and Malacaria, P., Full Abstraction for PCF. *Inf. and Comp.* 163, 409–470, 2000.
- [5] Bednarczyk, M.A., Hereditary History Preserving Bisimulations or What is the Power of the Future Perfect in Program Logics. ICS PAS Report, April 1991. Available from <http://www.ipipan.gda.pl/~marek/papers.html>.
- [6] Aghasaryan, A., Fabre, E., Benveniste, A., Boubour, R., and Jard, C., Fault detection and diagnosis in distributed systems: an approach by partially stochastic Petri Nets. *Discrete event dynamic systems, theory and applications*, vol 8 No 2, 1998.
- [7] Berry, G., Modèles complètement adéquats et stables des λ -calculs typés. Thèse de Doctorat d'Etat, Université de Paris VII, 1979.
- [8] Brock, J. and Ackerman, W., Scenarios: A model of non-determinate computation. *LNCS* 107, 1981.
- [9] Bunge, M., and Fiore, M.P., Unique factorisation lifting functors and categories of linearly-controlled processes. *MSCS* 10(2), 2000.
- [10] Carboni, A., and Vitale, E.M., Regular and exact completions *Pure and Applied Algebra* 125, 79–116, 1998.
- [11] Cattani, G.L., Stark, I., and Winskel, G., Presheaf models for for pi-Calculus. *Proc. CTCS'97*, *LNCS* 1290, 1997.
- [12] Cattani, G.L., and Winskel, G., Profunctors, open maps and bisimulation. *MSCS*, 2005.
- [13] Coquand, T., Gunter, C.A., and Winskel, G., DI-Domains as a Model of Polymorphism. *MFPS'97*, *LNCS* 298, 1987.
- [14] Crafa, S., Varacca, D., and Yoshida, N., Event Structure Semantics for the Internal pi-calculus. *CONCUR* 2007.
- [15] Crazzolaro, F., and Winskel, G., Events in security protocols. *ACM Conference on Computer and Communications Security*, 2001.
- [16] Curien, P.-L., Operads, clones, and distributive laws. Notes for an invited talk at Operads 2006, Strasbourg, Nov. 2006. Available from <http://www.pps.jussieu.fr/users/curien/>.
- [17] Curien, P.-L., and Faggian, C., L-Nets, Strategies and Proof-Nets. *CSL* 2005.
- [18] Danos, V., Feret, J., Fontana, W., Krivine, J., Scalable Simulation of Cellular Signaling Networks. *APLAS* 2007.
- [19] Emerson, E.A., and Sistla, A.P., Symmetry and model checking. *Formal Methods in System Design*, 1994.
- [20] Esparza, J., and Heljanko, K., Unfoldings: A Partial-Order Approach to Model Checking. *EATCS Monographs in Theoretical Computer Science*, 2008.
- [21] Fabre, E., Bayesian Networks of Dynamic Systems. Habilitation thesis, IRISA Rennes, 2007.
- [22] Fiore, M.P., Cattani, G.L., and Winskel, G., Weak Bisimulation and Open Maps. *LICS* 1999.
- [23] Fiore, M., Gambino, N., Hyland, J.M.E., and Winskel, G., The cartesian closed bicategory of generalised species of structures. *Journal of the London Math. Soc.*, 77 2, 203–220, 2007.
- [24] Ghica, D.R., and McCusker, G., Reasoning about Idealized Algol using regular languages. *ICALP* 2000, *LNCS* 1853, 2000.
- [25] Girard, J.-Y., Normal functors, power series and lambda calculus. *Ann. Pure Appl. Logic* 37, 1988.
- [26] Girard, J.-Y., Linear Logic. *Theoretical Computer Science*, 50:1102, 1987.
- [27] Girard, J.-Y., Towards a geometry of interaction. *Categories, Computer Science and Logic, Contemporary Mathematics*, AMS 92, 1989.
- [28] , R.B., Consistent Quantum Theory. *CUP*, 2003.
- [29] Glabbeek, R.J. van, On the expressiveness of higher dimensional automata. *EXPRESS* 2004, *ENTCS* 128(2), 2005.
- [30] Doghmi, S.F., Guttman, J.D., and Thayer, F.J., Searching for shapes in cryptographic protocols. *TACAS'07*, 2007.
- [31] Hayman, J., and Winskel, G., Independence and Concurrent Separation. *LICS* 2006.
- [32] Hayman, J., and Winskel, G., The unfolding of general Petri nets. Submitted 2008. Available from <http://www.cl.cam.ac.uk/users/gw104>.
- [33] Heiner, M., Gilbert, D., and Donaldson, R., Petri nets for Systems and Synthetic Biology. *LNCS* 5016, 2008.
- [34] Hoare, C.A.R., *Communicating Sequential Processes*. Prentice Hall, 1985.

- [35] Hyland, J.M.E., and Luke Ong, C-H., On Full Abstraction for PCF: I, II, and III. *Inf. Comput.* 163(2): 285-408, 2000.
- [36] Intel Corp., "Pentium Processor Users Manual," Vol.1.
- [37] Joyal, A., Une théorie combinatoire des séries formelles. *Advances in Mathematics* 42, 1981.
- [38] Joyal, A., Nielsen, M., and Winskel, G., Bisimulation from open maps. *LICS '93 special issue of Inf. and Comp.*, 127(2):164–185, 1996.
- [39] Kahn, G., The semantics of a simple language for parallel programming. *Information Processing*, vol. 74, 1974.
- [40] Lamport, L., Time, clocks and the ordering of events in a distributed system. *CACM* vol.21, 1978.
- [41] Lehmann D.J., Categories For Fixedpoint Semantics. Warwick University Theory of Computation Report No 15, 1976.
- [42] Nachbin, L., Topology and Order. Van Nostrand Mathematical Studies, 1965.
- [43] Pitts, A.M., Nominal Logic, A First Order Theory of Names and Binding. TACS2001 Special issue, *Inf. and Comp.* 186, 2003.
- [44] Petri, C. A., Kommunikation mit Automaten. Ph. D. Thesis. University of Bonn, 1962.
- [45] Martin, K., and Panangaden, P., A Domain of spacetime intervals for General relativity. In *Communications of Mathematical Physics*, Nov 2006.
- [46] McMillan, K.L., A technique of state space search based on unfolding. *Formal Methods in System Design* 6(1), 1995.
- [47] Milner, R., Communication and Concurrency. Prentice-Hall International, 1989.
- [48] Mislove, M.W., On Combining Probability and Nondeterminism. *ENTCS* 162: 261-265, 2006.
- [49] Moggi, E., Computational lambda-calculus and monads. *LICS* 1989.
- [50] Moore, S., and Greenfield, D., The Next Resource War: Computation vs. Communication. In the 10th International Workshop on System-Level Interconnect Prediction, 2008.
- [51] Nielsen, M., Krukow, K., and Sassone, V., A Bayesian Model for Event-based Trust. *ENTCS* 172: 499-521, 2007.
- [52] Nielsen, M., Plotkin, G.D., and Winskel, G., Petri nets, event structures and domains. *TCS*, 13(1):85–108, 1981.
- [53] Nygaard, M., Domain theory for concurrency. PhD Thesis, University of Aarhus, 2003.
- [54] Nygaard, M., and Winskel, G., Domain theory for concurrency. *TCS* 316: 153–190, 2004.
- [55] Hildebrandt, T., Panangaden, P., Winskel, G., A relational model of non-deterministic dataflow. *MSCS* 14(5): 613-649, 2004.
- [56] Plotkin, G.D., A Powerdomain Construction. *SIAM J. Comput.* 5(3): 452-487, 1976.
- [57] Plotkin, G. D., A Structural Approach to Operational Semantics. Technical Report DAIMI FN-19, Aarhus University, 1981.
- [58] Plotkin, G.D., A calculus of biochemical systems. In preparation, 2008.
- [59] Sassone, V., An axiomatization of the category of Petri nets computations. *MSCS* 8, 1998.
- [60] Saunders-Evans, L., and Winskel, G., Event structure spans for non-deterministic dataflow. *Proc. Express'06, ENTCS*, 2006.
- [61] Saunders-Evans, L., Events with persistence. PhD thesis, University of Cambridge Computer Laboratory, 2007.
- [62] Scott, D.S., Data types as lattices. *Lecture Notes in Mathematics* 499, 1975
- [63] Selinger, P., and Valiron, B., On a fully abstract model for a quantum linear functional language. *QPL 2006, Oxford, ENTCS* 210:123-137, 2008.
- [64] Sorkin, R., Geometry from order: causal sets. From http://www.einstein-online.info/en/spotlights/causal_sets/, 2006.
- [65] Stoy, J., Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory. MIT Press, 1981.
- [66] Turner, D., and Winskel, G., A nominal higher order process language and its semantics. In preparation, 2008.
- [67] Varacca, D., Two Denotational Models for Probabilistic Computation. PhD University of Aarhus, 2003.
- [68] Varacca, D., and Winskel, G., Distributing Probabililty over Nondeterminism. *MSCS* 16(1), 2006.
- [69] Varacca, D., Voelzer, H., and Winskel, G., Probabilistic event structures and domains. *TTCS* 358(2-3): 173-199, 2006.

- [70] Winskel, G., and Nielsen, M., Models for Concurrency. Handbook of Logic and the Foundations of Computer Science, vol. 4, pages 1-148, OUP, 1995.
- [71] Winskel, G., *Events in Computation*. PhD thesis, Univ. of Edinburgh, 1980. Available from <http://www.cl.cam.ac.uk/users/gw104>.
- [72] Winskel, G., Event structure semantics of CCS and related languages. ICALP 82, LNCS 140, 1982. Extended version available from <http://www.cl.cam.ac.uk/users/gw104>.
- [73] Winskel, G., A Note on Powerdomains and Modality. FCT, LNCS 158, 1983.
- [74] Winskel, G., A new definition of morphism on Petri Nets. STACS'84, LNCS 166, 1984.
- [75] Winskel, G., Event structures. LNCS 255, 1987.
- [76] Winskel, G., Linearity and nonlinearity in distributed computation. In the book 'Linear Logic in Computer Science,' CUP, 2004.
- [77] Winskel, G., Relations in concurrency. Invited talk, LICS'05, 2005.
- [78] Winskel, G., Event structures with symmetry. In the Plotkin Festschrift. ENTCS 172, 2007. See <http://www.cl.cam.ac.uk/users/gw104> for corrections.
- [79] Winskel, G., Symmetry and Concurrency. CALCO 2007.
- [80] Winskel, G., The symmetry of stability. In preparation, 2008.
- [81] Khomenko, V., Koutny, M., and Yakovlev, A., Logic Synthesis for Asynchronous Circuits Based on STG Unfoldings and Incremental SAT, Fundamenta Informaticae 70 (1-2), 2006.

A. APPENDIX: PSEUDO EQUIVALENCE

Assume a category with pullbacks. A pair of maps $l, r : S \rightarrow E$ forms a *pseudo equivalence* provided it is:

Reflexive: there is a map ρ such that

$$\begin{array}{ccccc} & & E & & \\ & \swarrow \text{id}_E & \vdots \rho & \searrow \text{id}_E & \\ E & \xleftarrow{l} & S & \xrightarrow{r} & E \end{array}$$

commutes;

Symmetric: there is a map σ such that

$$\begin{array}{ccccc} & & S & & \\ & \swarrow r & \vdots \sigma & \searrow l & \\ E & \xleftarrow{l} & S & \xrightarrow{r} & E \end{array}$$

commutes;

Transitive: there is a map τ such that

$$\begin{array}{ccccc} & & P & & \\ & \swarrow & \vdots \tau & \searrow & \\ & S & & S & \\ & \swarrow l & \searrow r & \swarrow l & \searrow r \\ E & & E & & E \end{array}$$

commutes.

When l, r are *jointly monic* (i.e. for all maps $x, y : D \rightarrow S$, if $lx = ly$ and $rx = ry$, then $x = y$) they form an *equivalence relation*.