Mathematical Structures in Computer Science

http://journals.cambridge.org/MSC

Additional services for ***Mathematical Structures in Computer Science:***

Email alerts: Click here
Subscriptions: Click here
Commercial reprints: Click here
Terms of use : Click here



Sheaf semantics for concurrent interacting objects

Joseph A. Goguen

Mathematical Structures in Computer Science / Volume 2 / Issue 02 / June 1992, pp 159 - 191
DOI: 10.1017/S0960129500001420, Published online: 04 March 2009

Link to this article: http://journals.cambridge.org/abstract_S0960129500001420

How to cite this article:
Joseph A. Goguen (1992). Sheaf semantics for concurrent interacting objects. Mathematical Structures in Computer Science, 2, pp 159-191 doi:10.1017/S0960129500001420

Request Permissions : Click here

CAMBRIDGE JOURNALS

Downloaded from http://journals.cambridge.org/MSC, IP address: 147.188.128.74 on 23 Mar 2015

Math. Struct. in Comp. Science (1992), *vol.* 2, *pp.* 159–191

Sheaf semantics for concurrent interacting objects

JOSEPH A. GOGUEN †

Programming Research Group, Oxford University

Received 15 January 1991; revised 12 December 1991

This paper uses concepts from sheaf theory to explain phenomena in concurrent systems, including object, inheritance, deadlock, and non-interference, as used in computer security. The approach is very general, and applies not only to concurrent object oriented systems, but also to systems of differential equations, electrical circuits, hardware description languages, and much more. Time can be discrete or continuous, linear or branching, and distribution is allowed over space as well as time. Concepts from category theory help to achieve this generality: objects are modelled by sheaves; inheritance by sheaf morphisms; systems by diagrams; and interconnections by diagrams of diagrams. In addition, behaviour is given by limit, and the result of interconnection by colimit. The approach is illustrated with many examples, including a semantics for a simple concurrent object-based programming language.

1. Introduction

Many popular formalisms for concurrent systems are *syntactic* (or "formal") in the sense that they represent systems by expressions, and then reason about systems by manipulating the corresponding expressions. For example, Milner's CCS (Milner 1980), Hoare's CSP (Hoare 1985) and Bergstra's ACP (Bergstra and Klop 1984) provide *process algebras*, which represent systems by expressions in which the primitives for process combination are implicitly defined by sets of equations; a quite different formal approach to concurrency is Girard's linear logic (Girard 1986).

What we call *semantic*, or *model theoretic*, approaches, provide complete sets of possible behaviours for systems. Such approaches have received less attention than syntactic approaches, but are important as standards against which to test the soundness and completeness of syntactic approaches, and also for defining basic general concepts, such as deadlock and information flow. Moreover, they are closer to our physical intuition, can often describe examples in simple and natural ways, and integrate easily with such additional considerations as data structure, objects and constraints. Trace models, as used in CSP (Hoare 1985) and other process algebras, are a prototypical example. From this point of view, Petri nets (Reisig 1986), (labelled) transition systems (Winskel 1990), and synchronisation trees (Milner 1980) can also be seen as syntactic.

Actually, things are not quite so simple, because the approaches that we have lumped

† Also with SRI International, Menlo Park CA 94025. The research reported in this paper has been supported in part by grants from the Science and Engineering Research Council, and the Fujitsu Corporation.

together as "syntactic" really have varying degrees of semantics. For example, transition systems and synchronisation trees have been used as semantics for CCS and CSP; also, CSP has a "preferred" model, based on failures and refusals (Hoare 1985). Petri nets have been used as models for linear logic (e.g., Meseguer & Montanari 1988), and set theoretic models have been given for Hewitt's actor approach (Agha 1986). Moreover, CCS expressions have been used as models for temporal logic. One person's syntax is another person's semantics.

This paper proposes a new model theoretic approach to concurrency based on sheaves. Sheaf theory developed in mathematics for studying relationships between local and global phenomena, and has been applied in algebraic geometry, differential geometry, analysis, and even logic. It has also been given an abstract form using category theory (Grothendieck 1963, Gray 1965), which among other things provides some general results about limits that are used in this paper. From the point of view of concurrency theory, it seems suggestive to think of sheaves as a generalisation of trace models. Sheaves handle real time systems, and variation over space as well as over time, either discrete or continuous, in fact over any topological space, in a very natural way. The definition of sheaf involves a generalisation of the prefix closure condition for traces, and introduces an important new idea, called the *sheaf condition.*

Our main motivation for using sheaf theory in Computing Science is the desire to give semantics for "new generation" systems, such as object-based concurrent information systems, programming languages, and operating systems. This approach allows the integration of concurrency with objects, data abstraction, information hiding, etc., and also helps to illuminate the notion of inheritance. Having such a semantic approach can be a significant help when designing a new language; for example, a sheaf theoretic approach helped with the design of FUNNEL, a hardware description language (Stavridou *et al.* 1991).

The approach is *declarative* and *constraint based*, and does not require any distinction between inputs and outputs. This makes it easy to treat applications, such as constraint logic programming and electrical circuits, that are essentially *relational*, in the sense that they involve finding solutions for a number of simultaneous relations (which are not just functions); in particular, there may be more than one solution, or no solution. It seems possible that functional input/output thinking has held back progress in this area, by making it more difficult to treat systems that involve the interdependent origination of their observed behaviour. Another goal has been to obtain a general model theoretic framework, within which general concepts such as deadlock and non-interference can be defined independent of formalism, and within which a wide variety of approaches can be compared. It may be worth emphasising that true concurrency can be modelled, and that we are not at all dependent upon interleaving. In particular, Lilius (1991) has shown how to model Petri nets in our sheaf theoretic framework.

This paper builds on a much earlier paper (Goguen 1975) which used sheaf theory as part of a research programme on "Categorical General Systems Theory" (Goguen 1971, 1973, Goguen and Ginali 1978). My interest in this area was revived by the desire to give a semantics for FOOPS (a Functional Object Oriented Programming System) (Goguen and Meseguer 1987, Goguen and Wolfram 1991) and for the Rewrite Rule Machine, a

multi-grain hierarchical massively parallel graph rewriting machine (see Goguen *et al.* 1989 and Goguen 1991b).

The main points made in this paper about the relationship between sheaves and objects are the following:

— objects give rise to sheaves;

— inheritance relations correspond to sheaf morphisms;

— systems are diagrams of sheaves; and

— the behaviour of a system is given by the limit of its diagram.

Here "diagram" and "limit" are to be understood in the sense of category theory. We also treat the interconnection of systems using colimits. Although this paper tries to give some philosophical motivation for these points, it does not attempt a serious philosophical treatment.

This paper has more definitions and examples than results. But the variety of examples may be surprising, and some of them illustrate a simple new approach to an important application area, the semantics of distributed concurrent (possibly object oriented) systems. Some of the definitions may also be surprising for their generality, including a notion of security that generalises the Goguen-Meseguer non-interference approach (Goguen and Meseguer 1982, 1984); indeed, this seems to be the most general such definition in the literature, and it applies, for example, to the security of real time distributed concurrent (possibly object oriented) databases. A definition of deadlock is also given, which again seems more general than anything in the literature.

Category theory has by now been used in many studies of concurrency; for example, see Winskel (1990), Meseguer & Montanari (1988), and Ferrari (1990). But as far as I know, only Monteiro and Pereira (1986) have previously studied concurrency using sheaves; however, their approach does not seem to be closely related to the present paper.

1.1. *Prerequisites and Notation*

Basic category theory and some intuition for concurrency are needed to read this paper. The former can be found in many places, including Mac Lane (1971), which requires some mathematical sophistication, and Goldblatt (1979), which may be especially recommended because it discusses sheaves, though in a different formulation from ours, and because it begins rather gently. An introduction to category theory for computing scientists is Barr and Wells (1990). Some underlying intuitions for basic categorical concepts are given in Goguen (1991a), and an overview of the Computing Science categorical literature is given in Pierce (1990). A basic introduction to concurrency using CSP is Hoare (1985).

We will use semicolon (";") to denote composition of functions, so that $(f;g)(x) = g(f(x))$; more generally, we let semicolon denote composition in any category, and we let 1_A denote the identity morphism at the object A. In set theory, we let ω denote the set of natural numbers, $\{0, 1, 2, ...\}$, and we let $\#S$ denote the cardinality of a set S.

2. Sheaves and Objects

Let us begin by asking how it is that we come to know about ordinary everyday objects. Consider a comfortable leather armchair C. Perhaps you see it briefly through a doorway, with only its left profile visible. Perhaps you later see it from the front. Maybe you eventually sit in it and notice its leathery smell, but perhaps you never see its back or bottom, and most likely you have no idea how it is stuffed. In fact, what you have is a collection of observations of certain attributes over certain regions of space-time; you never "have" the object as a whole. We can formalise such observations as functions $f : U \to A$ from some domain (of space-time) to some set A of attributes. If attributes $A_1, ..., A_n$ are observed, then we have $A = A_1 \times ... \times A_n$. Note that this is a *semantic* approach, based on direct observation of behaviour, rather than a *syntactic* approach, based on some sort of description or abstraction of behaviour.

The possible domains U are partially ordered by inclusion, and typically are closed under finite intersections and arbitrary unions, i.e., they form what is called a topological space. But for most of this paper, the following assumption about domains for observation is sufficient:

Definition 2.1. A **base** for observation is a family of sets partially ordered by inclusion.

The following examples describe some typical bases for observations that are of interest in Computing Science.

Example 2.2. For discrete time systems, the base consisting of intervals of natural numbers starting from an initial time 0 is often appropriate. Intuitively, 0 might represent the time when the system was created, or first observed, and the various intervals starting from 0 might represent periods of continuous observation of the system. Thus,

$$\mathscr{I}_0(\omega) = \{\emptyset, \{0\}, \{0,1\}, \{0,1,2\}, ...\} \cup \{\omega\},$$

where the set ω of all natural numbers is the domain for observations over a complete (infinite) life cycle. The notation $[n] = \{0, ..., n-1\}$ for $n > 0$, and $[0] = \emptyset$, is often convenient when using this base.

Example 2.3. Let $\mathscr{I}_0^f(\omega) = \mathscr{I}_0(\omega) - \{\omega\}$; i.e., the *finite* intervals starting from 0.

Example 2.4. The base consisting of semi-open intervals of real numbers starting at time 0 is often appropriate for real time systems,

$$\mathscr{I}_0(\mathscr{R}^+) = \{[0,r) \mid r \geq 0\} \cup \{\mathscr{R}^+\},$$

where $\mathscr{R}^+$ denotes the non-negative real numbers.

Example 2.5. A base consisting of certain unions of rectangles may be useful for describing the behaviour of certain systems distributed over one dimension in space and one dimension in time. In this example, $\mathscr{I}$ is the set of all subsets $U \subseteq \omega \times \omega$ that satisfy the following two conditions:

1 For $t \in \omega$, let $h(U,t) = \#\{h \mid (t,h) \in U\}$; then for each $t \in \omega$, we require that $h(U,t)$ is finite, and that $\{h \mid (t,h) \in U\} = [h(U,t)]$.
2 Given $t, t' \in \omega$ such that $|t - t'| = 1$, then $|h(U,t) - h(U,t')| \leq 1$.

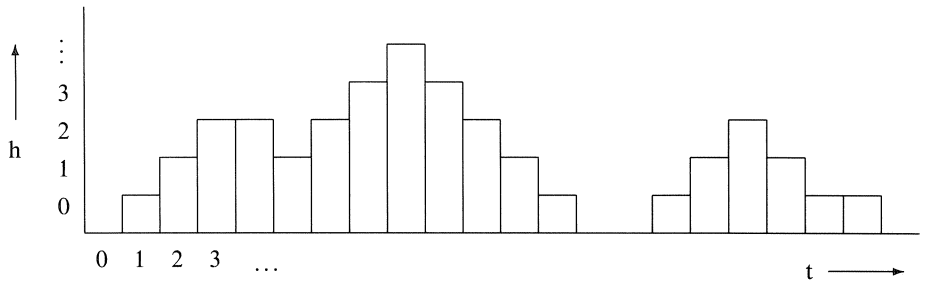


Fig. 1. An Open Set

These conditions say that each set in the base is a union of 1-by-n rectangles (for $n \geq 0$), such that the heights of two adjacent rectangles always differs by at most one; see Figure 1. We leave it to the reader to check that this base actually is a topological space.

The bases of Examples 2.2 and 2.4 are topological spaces, while the base of Example 2.3 is not, because it is not closed under arbitrary unions.

Example 2.6. If P is a preordered set (i.e., a set with a given relation $\leq$ that is reflexive and transitive), then the downclosed sets of P form a topological space, where a set $D \subseteq P$ is **downclosed** iff $d \in D$ and $d' \leq d$ imply $d' \in D$. This generalises Example 2.2 above. Let us denote this topology by $\mathscr{T}(P)$. An intuition that often works for this topology is that $d' \leq d$ means that "if you are now at d then you must have been at d'"; for example, this works when D is time.

For example, this construction can be used to define bases that reflect branching (non-linear) time, as used in temporal logic.

Any base can be embedded in a unique least topological space, by closing under finite intersections and arbitrary unions. Hence, little is lost by assuming that the base is a topological space. Also, it will be important for future discussions to note that any base can be regarded as a subcategory of the category $\mathscr{S}et$ of sets, with the domains as its objects, and their inclusions as its morphisms.

Given any inclusion $U \subseteq V$ and any observation $f : V \to A$, we can form the *restriction* of f to U, denoted $f \restriction U$; this is a function $U \to A$ which has the same values as f, but is only defined on the points in U.

Given an object (in the intuitive sense) $\mathcal{O}$, we let $\mathcal{O}(U)$ denote its set of observations over U, and given an inclusion $i : U \hookrightarrow V$, we let $\mathcal{O}(i) : \mathcal{O}(V) \to \mathcal{O}(U)$ denote the function that maps each $f : V \to A$ to its restriction $f \restriction U$. Then $\mathcal{O}$ satisfies the following two equations:

$$\mathcal{O}(i;j) = \mathcal{O}(j);\mathcal{O}(i)$$

where $j : V \hookrightarrow W$ is another inclusion and where ";" indicates the composition of functions; and

$$\mathcal{O}(1_U) = 1_{\mathcal{O}(U)}$$

where 1_U denotes the identity function on U. These two equations say that $\mathcal{O}$ is a

contravariant functor on the base $\mathscr{T}$ viewed as a category with inclusions as its morphisms. Since each $\mathcal{O}(U)$ is a set, we can view $\mathcal{O}$ as a set-valued functor $\mathcal{O} : \mathscr{T}^{op} \to \mathscr{S}et$, where $\mathscr{T}^{op}$ denotes the opposite category to $\mathscr{T}$. Although set-valued functors are often especially convenient (e.g., because one can apply the Yoneda lemma), the above formulation also suggests a natural generalisation, in which attributes might have another structure, such as that of a vector space, topological space, or Banach space. This is described in the following:

Definition 2.7. A **presheaf** is a functor $\mathcal{O} : \mathscr{T}^{op} \to \mathscr{C}$, where $\mathscr{T}$ is called its **base category** and $\mathscr{C}$ its **structure category**. If $i : U \hookrightarrow V$ is in $\mathscr{T}$, then $\mathcal{O}(i) : \mathcal{O}(V) \to \mathcal{O}(U)$ is called the **restriction** morphism induced by i. Given a preobject $\mathcal{O}$, the notation $f \restriction U$ is usually clearer than $\mathcal{O}(i)(f)$, where $f \in \mathcal{O}(V)$ and $i : U \hookrightarrow V$, and this notation is much used in the following.

This definition encapsulates the insight that observations are closed under restriction; this is a generalisation of the *prefix closure* condition of trace models. The fact that this condition can be formulated as functoriality allows the natural application of a number of basic results in category theory.

In all of our examples, the elements of each $\mathcal{O}(U)$ are functions, and in most of our examples, the structure category $\mathscr{C}$ is the category $\mathscr{S}et$ of sets. Section 2.2 below describes some modest (but somewhat more sophisticated) assumptions that allow us to handle other structures.

This approach can handle distribution over space, which is useful for studying multiprocessor computer systems, and it can also take account of continuity, linearity, or any other special structure that observations may have, by appropriate choice of $\mathscr{T}$ and $\mathscr{C}$. In fact, all of the presheaves in the examples of this paper arise in the following way:

Definition 2.8. A **preobject** $\mathcal{O}$ over a base $\mathscr{T}$ with **attribute object** $A = A_1 \times ... \times A_n$ is a presheaf of the form

$$\mathcal{O}(U) = \{h : U \to A_1 \times ... \times A_n \mid K(h)\},$$

where the morphisms $\mathcal{O}(i)$ are restriction maps, and the relation K expresses some property of functions, embodying the "laws" that $\mathcal{O}$ satisfies; the elements of A may be called **states, attributes** or **events**, depending on the context.

For example, K might arise from the laws defining a logic gate, an electronic device, or the distribution of heat on a sphere.

When $\mathscr{T} = \mathscr{I}_0(\omega)$, the elements of $\mathcal{O}(\{0, ..., i\})$ can be thought of as *traces* of the behaviour of $\mathcal{O}$ over U. (The states in these traces may be sets of more elementary states, or may have some other complex structure.) Intuitively, if $U \subseteq V$ and $h \in \mathcal{O}(U)$ and $h' \in \mathcal{O}(V)$, then $h = h' \restriction U$ means that the states in h' can evolve from those in h.

Sheaf theory suggests that the following additional condition may be quite fundamental for some applications:

Definition 2.9. An **object** is a preobject $\mathcal{O}$ such that its base $\mathscr{T}$ is a topological space, and such that given $U_i \in \mathscr{T}$ and $f_i \in \mathcal{O}(U_i)$ for all $i \in I$ such that $U = \bigcup_{i \in I} U_i$ and such that $f_i \restriction (U_i \cap U_j) = f_j \restriction (U_i \cap U_j)$ for all $i, j \in I$, then there is a unique $f \in \mathcal{O}(U)$ such that

$f \restriction U_i = f_i$ for all $i \in I$. This condition is called the **sheaf condition**. If the index set I is restricted to be finite, then the corresponding condition is called the **finite sheaf condition**.

The sheaf condition says that any set of pairwise consistent local observations can be "fused" together into a unique observation over the union of their domains. A deterministic input/output system has a subset of attributes (its inputs and/or states) such that knowing their values at some point in time uniquely determines the values of the other attributes (the outputs). However, the sheaf condition does not exclude non-determinism, in either the weak sense that there is more than one $f' \in \mathcal{O}(V)$ such that $f' \restriction U = f$ for some given $f \in \mathcal{O}(U)$ with $U \subseteq V$, or in the stronger sense that there is no subset of the attributes which is sufficient to determine all of each $f : U \to A$ in $\mathcal{O}(U)$. Note that non-determinism is *not* being modelled by sets of values in this discussion, although that could certainly be done for some applications if desired. Also, notice that if $\mathcal{T}$ is a compact topological space, then the finite and infinite sheaf conditions are equivalent; however, compactness does not seem to be very common in the applications.

The sheaf condition appears to be satisfied by the behaviours of all naturally arising systems from Computing Science. This "Sheaf Hypothesis" is similar to the so-called Church (or Church-Turing) thesis, that all intuitively computable functions are computable in the precise sense of (say) Turing machines. Indeed, the claim of Scott and others that computable functions are continuous can be seen as a special case of our claim (see Section 3.3 for some related discussion). However, the sheaf condition is *not* satisfied by certain *properties* of systems, when they are expressed as presheaves. For example, we will see in Example 2.16 that the property of *fairness* gives a presheaf that is not a sheaf. However, the finite sheaf condition *is* satisfied by such examples, and seems to be satisfied by all naturally arising *properties* of systems in Computing Science. Such property presheaves also satisfy the so-called *glueing condition*, which is the sheaf condition with "there exists a unique" replaced by just "there exists a".

Summing up the above discussion, our first main principle, which has much support from mathematical experience with many different kinds of geometrical object, as well as strong intuitions from concurrency theory in Computing Science, is that

OBJECTS GIVE RISE TO SHEAVES.

Sheaves can be thought of as a kind of "phase space" for objects. This sense of "object" might better be called an "object template" or "ideal object," because it describes all possible behaviours. When an object is actually part of a system, it may not exhibit all of its potential behaviours, because some of these may be inconsistent with constraints imposed by other objects. Also, objects in the present sense do not have unique identifiers; we will see later on that when part of a system, an object acquires a unique identifier (it will be the node that the object labels in the diagram of the system).

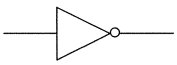
2.1. *Some Examples*

This subsection gives a number of examples illustrating the above definitions; some come from Computing Science, while others come from Electrical Engineering. Some more sophisticated examples are given in Section 4.

Example 2.10. (Inverter): For the discrete time case, $\mathscr{T} = \mathscr{I}_0(\omega)$, and when there is no delay, the behaviour of an inverter is described by the following sheaf,

$$\mathscr{O}(I) = \{f : I \to 2 \times 2 \mid i \in I \text{ and } f(i) = \langle t, t' \rangle \text{ imply } t' = \neg t\},$$

where $I \in \mathscr{T}$, $2 = \{0, 1\}$, and $\neg$ denotes negation (i.e., $\neg t = 1 - t$). The traditional picture is



If we impose a unit delay, we get instead the following:

$$\mathscr{O}(I) = \{f : I \to 2 \times 2 \mid i, i+1 \in I \text{ and } f(i) = \langle t, t' \rangle \text{ imply } f(i+1) = \langle t'', \neg t \rangle\}.$$

Example 2.11. (Gates): More generally, if $L(f_1, ..., f_n)$ is any Boolean-valued function of Boolean variables $f_1, \ldots, f_n$, then we can form a discrete time "L-gate" as follows, again for $I \in \mathscr{T} = \mathscr{I}_0(\omega)$:

$$\mathscr{O}(I) = \{f : I \to 2^{n+1} \mid i \in I \text{ implies } f_{n+1}(i) = L(f_1(i), ..., f_n(i))\},$$

where $f(i) = \langle f_1(i), ..., f_{n+1}(i) \rangle$. For example, an inverter arises by taking $L(t) = \neg t$, an AND-gate arises from $L(t, t') = t \wedge t'$, and an OR-gate arises from $L(t, t') = t \vee t'$.
For the unit delay case, we instead define, for $I \in \mathscr{I}_0(\omega)$,

$$\mathscr{O}(I) = \{f : I \to 2^{n+1} \mid i, i+1 \in I \text{ implies } f_{n+1}(i+1) = L(f_1(i), ..., f_n(i))\}.$$

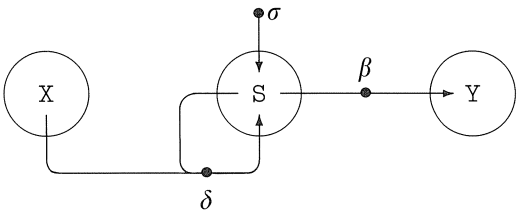
For example, a unit delayer arises by taking $L(t) = t$.

Example 2.12. (Automata): More generally still, we can use automata to model dynamic digital circuits whose behaviour depends on state as well as on input. The "ADJ diagram[1]" (i.e., algebraic signature[2]) given below describes deterministic (not necessarily finite state) automata with initial state σ, transition function δ, and output function β. Such an automaton A consists of three sets, A_X, A_S, and A_Y whose elements are called *inputs*, *states* and *outputs* respectively, plus an element $A_\sigma \in A_S$ and functions $A_\beta : A_S \to A_Y$ and $A_\delta : A_X \times A_S \to A_S$, called the *output* and *state transition functions*, respectively. When

[1] This name was suggested by Cliff Jones for a kind of diagram introduced by Goguen, Thatcher, and Wagner (Goguen *et al.* 1976) in their study of abstract data types. (The reason for the name "ADJ" is that the set {Goguen, Thatcher, Wagner, Wright} called itself ADJ at that time.)

[2] A **many sorted signature** Σ consists of a set S of sorts and a family of sets $\Sigma_{w,s}$ of operation symbols, one for each $w \in S^*$ and $s \in S$; we say that $\sigma \in \Sigma_{w,s}$ has **arity** w and **value sort** s. An operation symbol $\sigma \in \Sigma_{w,s}$ is interpreted as an operation $A_\sigma : A_w \to A_s$, where $A_w = A_{s_1} \times ... \times A_{s_n}$ when $w = s_1 ... s_n$. See Goguen *et al.* (1976) for more details.

just one automaton is under consideration and there is no danger of confusion, we will write X, S, Y for A_X, A_S, A_Y, and σ, β, δ for $A_\sigma, A_\beta, A_\delta$ respectively.



Here again time is discrete, so we let $\mathscr{T} = \mathscr{I}_0(\omega)$. The sheaf for this system is given, for $I \in \mathscr{T}$, by

$$\mathscr{O}(I) = \left\{ f : I \to X \times S \times Y \;\middle|\; \begin{array}{l} f_2(0) = \sigma \\ f_3(i) = \beta(f_2(i)) \text{ for all } i \in I \\ f_2(i+1) = \delta(f_1(i+1), f_2(i)) \text{ for all } i, i+1 \in I \end{array} \right\}$$

where f_1, f_2, f_3, respectively, denote the X, S, Y components of f. Automata can be used, for example, to define digital devices that have internal states, such as flip-flops.

Example 2.13. (Non-deterministic Automata): The only difference between deterministic and non-deterministic automata is that the transition function has the form $\delta : X \times S \to 2^S$ instead of $\delta : X \times S \to S$; that is, it returns a *set* of states rather than a single state. To get the sheaf for such a system, we replace the last line of the definition of $\mathscr{O}(I)$ in Example 2.12 above by

$$f_2(i+1) \in \delta(f_1(i+1), f_2(i)) \quad \text{for all} \quad i, i+1 \in I$$

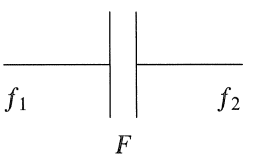
Because $\delta(f_1(i+1), f_2(i))$ can be empty, it may happen that there is no $h' \in \mathscr{O}([i+1])$ such that $h' \restriction [i] = h$ for some given $h \in \mathscr{O}([i])$; even for deterministic automata, it can happen that there is more than one such h', but for non-deterministic automata, there can be multiple h' having the same input components.

We can also model electrical components, using the real time base $\mathscr{I}(\mathscr{R}^+)$. For example:

Example 2.14. (Capacitor): Here, for $I \in \mathscr{I}(\mathscr{R}^+)$,

$$\mathscr{O}(I) = \left\{ f : I \to \mathscr{R}^3 \;\middle|\; \begin{array}{l} f \text{ is } C^\infty \text{ on } I, \text{ and} \\ f_3 = F \bullet \frac{d}{dt}(f_1 - f_2) \end{array} \right\},$$

where F is the capacitance, f_3 is the current, f_1, f_2 are voltages, and where a C^∞ function has continuous derivatives of every order on its domain. The following is the traditional picture:



Classical electrical engineering, which encompasses only linear devices, such as capacitors, inductors, and resistors, leads to systems of first order linear differential equations, of the general form

$$a_1\frac{dx_1}{dt} + a_2\frac{dx_2}{dt} + \cdots + a_n\frac{dx_n}{dt} + b_1x_1 + b_2x_2 + \cdots + b_nx_n = c\ .$$

In this setting, we could choose to exploit a more sophisticated structure category: because the differential equations involved are linear, the spaces of solutions are vector spaces; and because they also have a topological structure, we could let $\mathscr{C}$ be the category of topological vector spaces; an even better choice would be the category of Banach spaces (See Yosida (1968) for information on functional analysis, including Banach spaces, and Section 2.2 below for an approach to handling such additional structure.)

In much the same way, we could look at an object of solutions to some partial differential equation on a smooth manifold, for example, describing the flow of heat on the surface of a sphere. This would involve a 3-dimensional space[3], say embedded in 4-dimensional Euclidean space. See Goguen (1975) for further examples along similar lines. We now return to discrete time.

Example 2.15. (Networks): Networks are usually built from sites and links[4]. Clearly the sites can be regarded as objects, and for many purposes it is also convenient to regard the links as objects. For example, it is not possible to send arbitrary signals over a real link, because it has a certain bandwidth, certain conventions about how signals are represented, certain physical properties, etc. Thus, we represent a network by a graph whose nodes include both the sites and the links of the network, and whose edges represent a connection from a node to a link. For example, consider a network with four sites, labelled $P1, P2, P3, P4$, and four links, labelled $F1, F2, F3, F4$, as in the figure of Example 4.9 below. We can regard such a graph as defining a *preorder*, with $n > n'$ iff there is a link from n to n'. For example, $P1 > F1$. The downclosed sets then give a topology, in which an open neighbourhood N is a collection of sites and links such that if a site is in N, then all the links connected to it are also in N. The topologies of various SIMD and MIMD multi-processor architectures can be handled in this way.

Real networks are *dynamic* in the sense that some sites and/or links may sometimes be "down" and other times be "up." Some networks "dynamically reconfigure" themselves. It might seem that such networks cannot be modelled with the fixed topologies discussed above. However, they can be, by fixing in advance all potential communication links, and including "up/down" as part of the state of each node. For example, if the potential sites are $\{S_i \mid i \in \omega\}$, and if there is a potential link $L_{i,j}$ between each pair (i, j) of sites, then we let

$$N = \{S_i \mid i \in \omega\} \cup \{L_{i,j} \mid i, j \in \omega\}$$

where $L_{i,j} < S_i, S_j$ for all $i, j \in \omega$. Similarly, for a "star" topology, we would let

[3] The third dimension is for time.

[4] The words "node" and "edge" would perhaps be clearer, but would conflict with the use of these words in connection with graphs later in this discussion.

$$N = \{S_i \mid i \in \omega\} \cup \{\star\}$$

where $\star$ denotes the central node, with $\star < S_i$ for each $i \in \omega$.

If P is a totally ordered set (i.e., for all $d, d' \in P$ either $d \leq d'$ or $d' \leq d$ (or both)), then the finite sheaf condition is automatically satisfied in the downclosed topology $\mathscr{T}(P)$. It follows from this that any presheaf over the base $\mathscr{I}_0^f(\omega)$ is a sheaf, because any non-trivial infinite union would have union ω, which has been excluded. However, not all presheaves over $\mathscr{I}_0(\omega)$ and $\mathscr{I}_0(R^+)$ are sheaves, because interesting phenomena can appear "at infinity" that do not appear in the finite approximations, as shown in the following:

Example 2.16. (Fair Scheduler): We use the base $\mathscr{I}_0(\omega)$, and consider a "fair scheduler" $\mathscr{F}$ for two events a, b. What we mean by fairness here is that if an a occurs, then eventually a b must occur, and if a b occurs, then eventually an a must occur. This means that

$$\mathscr{F}(\omega) = (a^+b^+ + b^+a^+)^\omega,$$

which is a concatenation of strings, the first consisting of some a's followed by some b's or else some b's followed by some a's, and so on forever. It is interesting now to notice that for each $n \in \omega$,

$$\mathscr{F}([n]) = \{a, b\}^n,$$

i.e., *any* behaviour is possible on each finite interval $\{0, ..., n-1\}$, including all a's, and all b's.
If we now suppose that $\mathscr{F}$ is a sheaf, then the sheaf condition implies that

$$\mathscr{F}(\omega) = \{a, b\}^\omega,$$

i.e., that *any* behaviour is also possible in the limit. But this contradicts the definition of $\mathscr{F}$; therefore $\mathscr{F}$ is *not* a sheaf. However, it does satisfies the finite sheaf and gluing conditions.

Example 2.17. (Stacks): We can give an implementation for a stack that is distributed in space, as well as in time, using the base of Example 2.5, by defining

$$\mathscr{O}(U) = \{f : U \to \omega \mid h \in h(U,t) \cap h(U,t+1) \Rightarrow f(t,h) = f(t+1,h)\}.$$

(It might not appear so at first, but this example does satisfy the infinite sheaf condition.)

2.2. *Structure Categories*

This subsection presents some assumptions on the structure category $\mathscr{C}$ that will make everything work just as well as in the set case. (This material can be skipped on a first reading, as it requires somewhat more sophisticated category theory.)

Because Definition 2.8 involves an attribute object $A = A_1 \times ... \times A_n$, we must assume that structure categories have finite products. In fact, we will see later that we need limits over all small (or at least finite) diagrams, i.e., we need (finite) *completeness*. All of the

proposed structure categories mentioned so far in fact satisfy these assumptions, so they are not very restrictive in practice.

Definition 2.18. A **structure category** is a (finitely) complete category $\mathscr{C}$ together with a functor $\mathscr{U} : \mathscr{C} \to \mathscr{S}et$, called its **forgetful functor**, that has a left adjoint $F : \mathscr{S}et \to \mathscr{C}$, called its **free functor**.

For example, if $\mathscr{C}$ is the category of vector spaces, then $\mathscr{U}$ gives the underlying set of a vector space, and F gives the free vector space using a given set as basis. In this connection, Definition 2.9 should be made more precise, by replacing "$u : U \to A_1 \times ... \times A_n$" by "$u : U \to \mathscr{U}(A_1 \times ... \times A_n)$".

Proposition 2.19. If $\mathscr{C}$ is a structure category, then its forgetful functor $\mathscr{U}$ preserves products, and more generally, all limits.

Proof. It is well known that right adjoints preserve limits; e.g., see Mac Lane (1971) or Goldblatt (1979). □

This means that the underlying set of a product object in $\mathscr{C}$ can be taken to be the product of the underlying sets of the component objects; for example, we can get a product of two vector spaces by giving a vector space structure to the product of their underlying sets. In the following section, we will see that limits give the behaviour of systems; hence they are quite basic to our approach.

Here is how the definition of object looks in the present general setting:

Definition 2.20. A preobject $\mathscr{O} : \mathscr{T}^{op} \to \mathscr{C}$, with structure category $\mathscr{C}$ and with base $\mathscr{T}$ a topological space, satisfies the **finite sheaf condition** iff the following is an equalizer diagram in $\mathscr{C}$,

$$\mathscr{O}(U \cup V) \longrightarrow \mathscr{O}(U) \times \mathscr{O}(V) \rightrightarrows \mathscr{O}(U \cap V)$$

for all $U, V \in \mathscr{T}$, where the first arrow is the tupling of the restrictions to U and V, and the next two are $\pi_1 ; \restriction (U \cap V)$ and $\pi_2 ; \restriction (U \cap V)$. $\mathscr{O}$ is a **sheaf** iff the following is an equalizer diagram in $\mathscr{C}$,

$$\mathscr{O}(U) \longrightarrow \Pi_j \, \mathscr{O}(U_j) \rightrightarrows \Pi_{j,j'} \, \mathscr{O}(U_j \cap U_{j'}).$$

whenever $U = \bigcup_j U_j$, where the first arrow is the tupling of the restrictions, and the next two are given as compositions of projections with restrictions. This is called the **sheaf condition**.

It seems worth noting at this point that every presheaf is contained in a least sheaf, called its *enveloping sheaf*; more precisely, there is a left adjoint to the inclusion functor of sheaves into presheaves, called the *sheafification* functor; see Tennison (1975) for some further details.

3. System, Behaviour and Interconnection

The previous section presented objects as coherent collections of possible observations. This section considers systems of such objects, and in particular, it considers how their joint behaviour is determined by the behaviours of their components. It also considers

relationships of inheritance between objects, and how to interconnect systems. A very general compositionality result is proved for these notions. This approach seems especially convenient for systems that are highly concurrent, such as electrical circuits (whose component objects are capacitors, resistors, inductors, etc.) and digital circuits (whose components are flip flops, inverters, AND gates, etc.).

3.1. *Systems*

This subsection argues that systems can be modelled by diagrams of sheaves. In order to consider systems as diagrams, we need an appropriate notion of morphism between objects; these must be able to express the kinds of relationships between objects that arise when they are connected together to form systems; we argue that such relationships can be regarded as instances of *inheritance*. We use preobject morphisms to express relationships between component objects in a system, defined as follows:

Definition 3.1. Given preobjects $\mathcal{O}$ and $\mathcal{O}'$ over the same base $\mathcal{T}$, a **morphism** $\varphi : \mathcal{O} \to \mathcal{O}'$ is a family $\varphi_U : \mathcal{O}(U) \to \mathcal{O}'(U)$ of maps, one for each $U \in \mathcal{T}$, such that for each $i : U \to V$ in $\mathcal{T}$, the diagram

$$\begin{array}{ccc} \mathcal{O}(V) & \xrightarrow{\varphi_V} & \mathcal{O}'(V) \\ \Big\downarrow{\scriptstyle \mathcal{O}(i)} & & \Big\downarrow{\scriptstyle \mathcal{O}'(i)} \\ \mathcal{O}(U) & \xrightarrow{\varphi_U} & \mathcal{O}'(U) \end{array}$$

commutes in $\mathscr{C}$. When $\mathcal{O}$ and $\mathcal{O}'$ are objects, we may also call φ an **object morphism** or a **sheaf morphism**. This gives rise to categories $\mathscr{P}re\mathcal{O}bj(\mathcal{T},\mathscr{C})$ and $\mathcal{O}bj(\mathcal{T},\mathscr{C})$ of preobjects and objects, respectively, over a given base $\mathcal{T}$ and structure category $\mathscr{C}$.

Actually, such a φ is just a natural transformation from the functor $\mathcal{O}$ to the functor $\mathcal{O}'$.

Lemma 3.2. Given preobjects $\mathcal{O}$ and $\mathcal{O}'$ with attribute objects A and A' respectively, and given $a : A \to A'$, then defining $a_U : [U \to A] \longrightarrow [U \to A']$ by $f \mapsto f;a$, i.e., by $a_U(f) = f;a$, gives a morphism iff $U \in \mathcal{T}$ and $f \in \mathcal{O}(U)$ imply $f;a \in \mathcal{O}'(U)$.

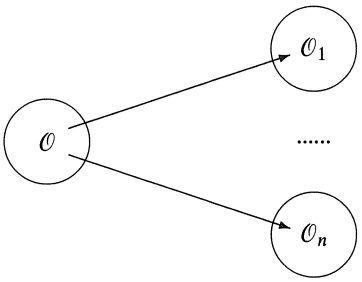
Definition 3.3. Given preobjects $\mathcal{O}$ and $\mathcal{O}'$ with attribute objects $A = \prod_{j\in J} A_j$ and $A' = \prod_{j\in J'} A_j$, respectively, with $J' \subseteq J$, if $a : A \to A'$ sending $\langle a_j \mid j \in J\rangle$ to $\langle a_j \mid j \in J'\rangle$ induces a morphism by the method of Lemma 3.2, then it is called a **projection morphism**.

For example, if $J = \{1,...,n\}$ and $J' = \{1,...,m\}$ with $m < n$, then we may think of the projection map $\mathcal{O} \to \mathcal{O}'$ as "forgetting" the attributes $A_{m+1},...,A_n$ of $\mathcal{O}$, or in a dual but perhaps more suggestive language, as expressing the *inheritance* by $\mathcal{O}$ of the attributes $A_1,...,A_m$ from $\mathcal{O}'$, with $A_{m+1},...,A_n$ added as its own local attributes. All of the morphisms that occur in the examples of this paper are projections. Non-projection morphisms describe a kind of generalised inheritance that may involve combining attributes and changing representation. These considerations motivate our second main principle, that

MORPHISMS REPRESENT INHERITANCE.

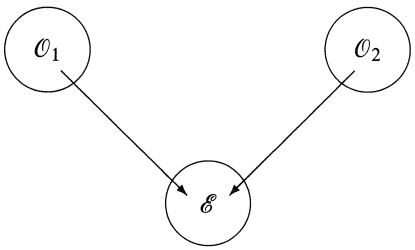
In particular, $\mathcal{O}$ inherits from $\mathcal{O}'$ iff $\varphi : \mathcal{O} \to \mathcal{O}'$ (the apparent reversal of direction here arises from the duality between "forgetting" and inheritance mentioned above).

In much the same way, *multiple inheritance* arises from multiple morphisms, as illustrated by the following diagram,



which we describe by saying that $\mathcal{O}$ inherits from $\mathcal{O}_1, ..., \mathcal{O}_n$.

Now let us return to the question of how objects form systems: in order for two objects to interact, they must each inherit something (some attributes and behaviour) from a third object; i.e., they must share a "common language" before they can speak to each other. This may be pictured as follows:

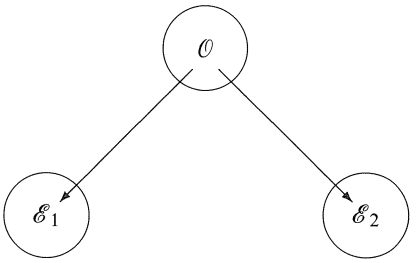


Let us call this a "valley" or "V" diagram. In many cases, the shared "language" $\mathcal{E}$ consists of *all possible* behaviours having states in some attribute object, that is, it has the form

$$\mathcal{E}(I) = \{f : I \to A \mid M(f)\}$$

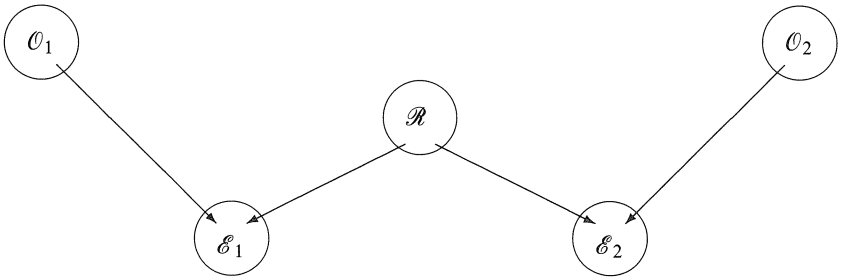
for each $I \in \mathcal{T}$, where the relation M may, for some application areas, express some minimal conditions (such as being linear, bounded, or continuously differentiable) and where A contains the states of some communication medium, such as the real numbers or an appropriate alphabet of events. It is common for $M(f)$ to hold for all f, and in general we expect M to be such that $J \subseteq I$ and $M(f)$ imply $M(f \restriction J)$; the latter condition is sufficient for $\mathcal{E}$ to be a functor. It may help to think of $\mathcal{E}$ as the general object of all possible event streams or traces.

Another common case is for some relationship to hold between two languages. This may be pictured as follows:



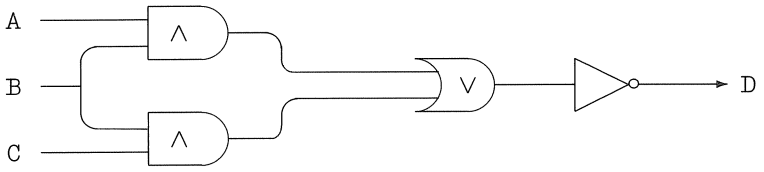
Let us call this a "peak" diagram. For example, an inverter is a relation between two objects of discrete time Boolean valued event streams. Similarly, AND and OR gates are relations between three such objects.

A more general way for two objects to communicate is for there to be a "boundary object" that can translate between their languages, or at least pass along some information, based on the existence of a relationship between the two languages. The system diagram for this is a "double valley" or "W" diagram, as follows:

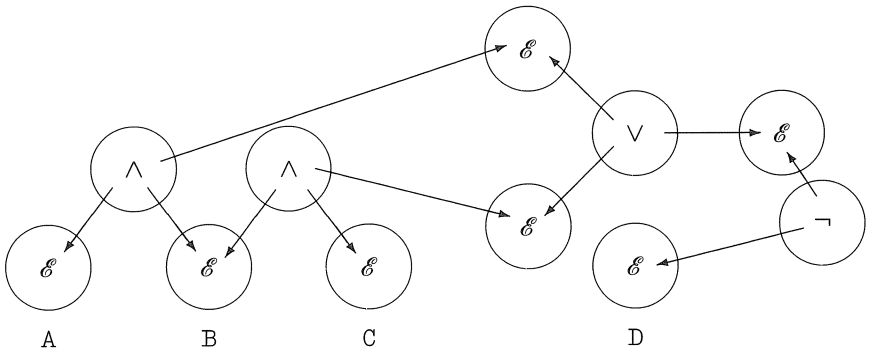


in which $\mathcal{O}_i$ has language $\mathcal{E}_i$ (for $i = 1, 2$) and $\mathcal{R}$ is a relation between these languages. For example, $\mathcal{E}_1$ might be all streams over discrete time and $\mathcal{E}_2$ all streams over continuous time, with $\mathcal{R} \to \mathcal{E}_1$ and $\mathcal{R} \to \mathcal{E}_2$ giving the translation of a discrete time interval $\{0, 1, ..., n\}$ to a continuous time interval $[0, n+1)$. Similarly, $\mathcal{R}$ might translate between two different clocks, thus allowing for asynchronous communication. For another example, $\mathcal{R}$ might represent the summarisation of complex data in $\mathcal{O}_1$ to simple statistics in $\mathcal{O}_2$. $\mathcal{O}_1$ and $\mathcal{O}_2$ can be thought of as two different perspectives on a more complex object $\mathcal{R}$ or as providing a form of non-strict inheritance. What is called "overriding" in object oriented programming can probably be handled in this way. (These ideas were inspired by work of Leigh Star (Star 1988) on "boundary objects" in the sociology of science.)

Example 3.4. Now let us consider the system that in traditional notation would be described by the following diagram,



where $\wedge, \vee$ label AND and OR gates, respectively, corresponds to the following diagram in the category of objects over $\mathscr{I}_0(\omega)$:



These considerations motivate our third main principle, that

SYSTEMS ARE DIAGRAMS.

In line with our attempt to keep the category theoretic prerequisites of this paper to a minimum, we may explain the above with the following:

Definition 3.5. A **system** $\mathscr{S}$ consists of a graph with nodes $n \in N$ labelled by (pre)objects $\mathscr{S}_n$ and with edges $e : n \to n'$ labelled by morphisms $\varphi_e : \mathscr{S}_n \to \mathscr{S}_{n'}$.

All of the (pre)objects and morphisms in $\mathscr{S}$ are assumed to have the same base and the same structure category. (We will give a more sophisticated definition of system later.)

3.2. *Behaviour*

This subsection argues that the behaviour of a system is given by the limit of its diagram. Let us begin by constructing an object that describes all the possible behaviours of a system. So, let us assume a system $\mathscr{S}$ with objects $\mathscr{S}_n$ for $n \in N$, and let us choose some fixed domain $I \in \mathscr{T}$. Then a possible behaviour of the system over I is a choice of one behaviour for each object, say $f_n : I \to A_n$ in $\mathscr{S}_n(I)$, such that this family f_n is mutually consistent, in the sense that for each morphism $\varphi_e : \mathscr{S}_n \to \mathscr{S}_{n'}$ we have $\varphi_e(f_n) = f_{n'}$. Let us call such a family $\{f_n \mid n \in N\}$ a **consistent net of points** in $\mathscr{S}$.

Thus, the object of behaviours of the system has, for each $I \in \mathscr{T}$,

$$\mathscr{L}(I) = \left\{ \{f_n \mid n \in N\} \;\middle|\; \begin{array}{l} f_n \in \mathscr{S}_n(I) \text{ and } \varphi_e : \mathscr{S}_n \to \mathscr{S}_{n'} \\ \text{imply } \varphi_e(f_n) = f_{n'} \end{array} \right\},$$

that is, it contains all of the consistent nets of points over I.

When the structure category $\mathscr{C}$ is $\mathscr{S}et$, it is well known (e.g., Mac Lane 1971, Goldblatt 1979), that $\mathscr{L}(I)$ is (a construction for) the *limit* of the diagram $\mathscr{S}(I)$ in which each node n is labelled by $\mathscr{S}_n(I)$ and each edge $e : n \to n'$ is labelled by $\varphi_e : \mathscr{S}_n(I) \to \mathscr{S}_{n'}(I)$. (One might say that the limit is "trying to make the diagram commute".)

We now draw upon another general result from category theory, showing that limits of diagrams of preobjects are computed "pointwise," i.e., we have the following, in which $\lim_n \mathscr{S}_n$ denotes the limit of a system $\mathscr{S}$ of objects $\mathscr{S}_n$, including their morphisms:

Proposition 3.6. Any diagram $\mathscr{S}$ of presheaves with values in a structure category $\mathscr{C}$ has a limit, and in fact, for each $U \in \mathscr{T}$,

$$\left(\lim_n \mathscr{S}_n\right)(U) = \lim_n(\mathscr{S}_n(U)).$$

Proof. This follows directly from a well known result about limits in functor categories that is proved, for example, in Mac Lane (1971) and Goldblatt (1979), using the assumption that the structure category has limits. □

This motivates our fourth main principle, that

BEHAVIOUR IS LIMIT.

Proposition 3.6 and our assumption that $\mathscr{C}$ has limits imply that every system has a behaviour object $\mathscr{L}$. But this does not mean that every system actually exhibits behaviour over every $I \in \mathscr{T}$; for example (assuming $\mathscr{C} = \mathscr{S}et$), it is possible that $\mathscr{L}(I) = \emptyset$ for some, or even for all, $I \in \mathscr{T}$.

To illustrate this principle, the **parallel composition** of two objects, $\mathscr{O}_1 \parallel \mathscr{O}_2$, is the sheaf given by their product, $\mathscr{O}_1 \times \mathscr{O}_2$, and their **synchronised parallel connection** is given by the limit of a valley diagram, where synchronising events occur in the bottom object.

The relationship between global and local behaviour that arises between a system and its component objects should be distinguished from the global/local relationship stated in the sheaf condition of Definition 2.9. The first concerns the behaviour of multi-object systems, through diagrams and their limits, while the second concerns "glueing together" behaviour over domains.

3.3. *Interconnection*

The principles that objects are sheaves, systems are diagrams, and behaviour is limit are all taken from some earlier work in categorical General System Theory (Goguen 1971, 1973, Goguen and Ginali 1978). Another principle from this work is that interconnecting systems corresponds to taking colimits in the category of systems, where sharing is indicated by inclusion maps from shared parts into the systems that share them. The

papers Goguen (1971, 1973) and Goguen and Ginali (1978) develop some very general results in this setting, including the so-called Interconnection and Behaviour Theorems, which are given below[5]. We apply this material to show that the behaviour of a sheaf at a limit point is the limit of its behaviours at approximating points. (This is more technical than most of the rest of the paper, and some readers may wish to skip it on a first reading.)

In order to have a category of systems, we first need to define morphisms of systems. It is convenient to do this in the general setting of diagrams over an arbitrary category[6] $\mathscr{S}$, as follows:

Definition 3.7. A **diagram** in a category $\mathscr{S}$ is a functor $D : B \to \mathscr{S}$ from some **base** category B. Given diagrams $D_0 : B_0 \to \mathscr{S}$ and $D_1 : B_1 \to \mathscr{S}$, a **morphism** from D_0 to D_1 consists of a functor $F : B_0 \to B_1$ and a natural transformation $\eta : F; D_1 \Rightarrow D_0$.
Given three diagrams, $D_i : B_i \to \mathscr{S}$ for $i = 0, 1, 2$, the **composition** of the morphisms $\langle F_1, \eta_1 \rangle : D_0 \to D_1$ and $\langle F_2, \eta_2 \rangle : D_1 \to D_2$ is a morphism $\langle F_1; F_2, (F_1 \circ \eta_2); \eta_1 \rangle : D_0 \to D_2$.

Here the operation ";" on natural transformations is their vertical composition, whereas "$\circ$" is their horizontal composition; e.g., see Mac Lane (1971). Although Definition 3.5 defined a system to be a kind of labelled graph, it is not difficult to see that any such labelled graph extends uniquely to a functor with source the category of paths in the given graph; and conversely, any functor can be considered a labelling of the underlying graph of its source category. So Definition 3.7 is consistent with Definition 3.5.

The direction of the arrows $\eta_n : D_1(F(n)) \to D_0(n)$ may seem counter-intuitive at first, but, for example, in the common case of an inclusion from a one node diagram into a system, it corresponds to the need for translating the language of an object in the system into that of the object in the one node diagram.

It is not difficult to check the following:

Fact 3.8. Composition of diagram morphisms is associative and has identities, so that we have a **category of diagrams in** $\mathscr{S}$, denoted $\mathbb{D}gm(\mathscr{S})$.

The following basic result guarantees that in the setting of this paper, where the category $\mathscr{S}$ of systems is complete, we can always interconnect systems. (The elegant proof is due to Răzvan Diaconescu.)

Theorem 3.9. (Interconnection Theorem) If $\mathscr{S}$ is (finitely) complete, then $\mathbb{D}gm(\mathscr{S})$ is also (finitely) cocomplete, i.e., has all (finite) colimits.

Proof. First, we define a functor $P : \mathbb{C}at^{op} \to \mathbb{C}at$ by

$$P(B) = [B \to \mathscr{S}]^{op} \text{ for any category } B, \text{and}$$
$$P(F) = [F \to \mathscr{S}]^{op} : [B_1 \to \mathscr{S}]^{op} \to [B_0 \to \mathscr{S}]^{op} \text{ for any functor } F : B_0 \to B_1.$$

Noting that P is a (strict) indexed category (in the sense of (Tarlecki *et al.* 1991)), and that $\mathbb{D}gm(\mathscr{S}) = \text{Flat}(P)$ (again, see (Tarlecki *et al.* 1991) for this notation), we can

[5] Actually, the results given here are somewhat more general than those in Goguen (1971, 1973) and Goguen and Ginali (1978), because the restriction to so-called interconnection morphisms has been removed.

[6] In particular, the objects of $\mathscr{S}$ need not be sheaves or presheaves.

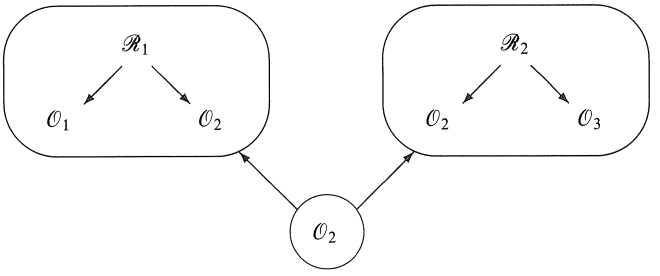


Fig. 2. The Composition of Two Relations

use Theorem 2 of (Tarlecki *et al.* 1991) to show the (finite) completeness of Flat(P) by checking the hypotheses of that theorem:

1 $\mathbb{C}at$ is cocomplete.

2 $[B \to \mathscr{S}]^{op}$ is (finitely) cocomplete for any category B because $\mathscr{S}$ is assumed (finitely) complete.

3 P is locally reversible because $P^{op} : [B_1 \to \mathscr{S}]^{op} \to [B_0 \to \mathscr{S}]^{op}$ has a left adjoint, because $P : [B_0 \to \mathscr{S}] \to [B_1 \to \mathscr{S}]$ has a right adjoint, because any factor $F : B_0 \to B_1$ has a right Kan extension along any functor $B_0 \to \mathscr{S}$ by the Kan Extension Theorem (see Theorem 1 on page 233 of (Mac Lane 1971)) because $\mathscr{S}$ is assumed complete.

□

Taking colimits in the category $\mathbb{D}gm(\mathscr{S})$ corresponds to interconnecting systems. For example, interconnecting two relations (represented by "peak" diagrams) over a common object gives a "double peak" diagram, whose behaviour (i.e., limit) gives the a relation which is the composition of the two given relations, as illustrated in Figure 2.

We may summarise the above discussion in the following principle, from Goguen (1971, 1973) and Goguen and Ginali (1978) (see also Goguen 1991a):

INTERCONNECTION IS COLIMIT.

We will soon need the following result about limits:

Proposition 3.10. Given a complete category $\mathscr{S}$, then "limit of" is a functor

$$\text{Lim} : \mathbb{D}gm(\mathscr{S})^{op} \longrightarrow \mathscr{S}$$

which is right adjoint to the functor

$$\lceil _ \rceil : \mathscr{S} \longrightarrow \mathbb{D}gm(\mathscr{S})^{op}$$

which sends each object C to the diagram consisting of just one node labelled C.

Proof. This is just a reformulation of the universal property of limits. In particular, note that the functor category $\mathbb{D}gm(\mathscr{S})[D, \lceil C \rceil]$ is the category of all cones $\eta : \lceil C \rceil \Rightarrow D$ over D. □

Intuition about systems suggests that one can calculate the behaviour of a system from

the behaviours of its components. This intuition is actually a precise theorem in the present formal setting, stated as follows:

Theorem 3.11. (Behaviour Theorem) Let $D : I \to \mathbb{D}gm(\mathscr{S})^{op}$ be a (small) diagram of diagrams (i.e., systems) over a complete category $\mathscr{S}$. Then (in a hopefully suggestive notation)

$$\mathrm{Lim}(\mathrm{colim}\ D) = \lim(D; Lim).$$

Proof. As noted in the proof of Proposition 2.19, right adjoints preserve limits; also, limits in $\mathbb{D}gm(\mathscr{S})^{op}$ are colimits in $\mathbb{D}gm(\mathscr{S})$. □

This result is a very general "*Compositionality Theorem*," in the sense of giving very general conditions under which behaviours of parts can be composed to give the behaviour of the whole.

Let us consider a special case, where I is the discrete base with just two nodes, say 0,1. Then a functor $D : I \to \mathbb{D}gm(\mathscr{S})^{op}$ consists of just two diagrams, D_0 and D_1. If we denote their colimit in $\mathbb{D}gm(\mathscr{S})$ by $D_0 + D_1$, then we have the following formula:

$$\lim(D_0 + D_1) = \lim(D_0) \times \lim(D_1).$$

It may be amusing, and perhaps surprising, that Theorem 3.11 can be used to prove that the behaviour of a sheaf at a "limit point" is the limit of the behaviours leading up to it. Given a presheaf $\mathscr{O} : \mathscr{T}^{op} \to \mathscr{C}$ with $\mathscr{C}$ complete, let $U = \bigcup_{j \in J} U_j$ in $\mathscr{T}$, for some index set J. Let I have the shape

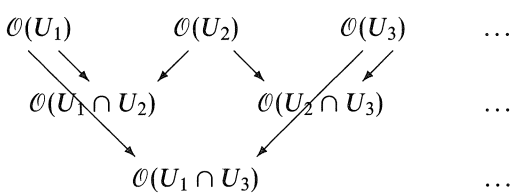
$$\{0\} \leftleftarrows \{1\}$$

We will now define a diagram D in $\mathbb{D}gm(\mathscr{T})$ with base I. Let D_0 be the diagram with base the set J and with $j \in J$ labelled by $\mathscr{O}(U_j)$; let D_1 be the diagram with base $J \times J$ and with $\langle j, j' \rangle$ labelled by $\mathscr{O}(U_j \cap U_{j'})$. Next, define $d : D_1 \to D_0$ as follows: on the bases, d_F is the projection $\langle j, j' \rangle \mapsto j$; and on the objects, $(d_\eta)_{\langle j,j' \rangle}$ is the restriction morphism $\mathscr{O}(U_j) \to \mathscr{O}(U_j \cap U_{j'})$. Define $d' : D_1 \to D_0$ similarly: on the bases, d'_F is the projection $\langle j, j' \rangle \mapsto j'$; and on the objects, $(d'_\eta)_{\langle j,j' \rangle}$ is the restriction morphism $\mathscr{O}(U_{j'}) \to \mathscr{O}(U_j \cap U_{j'})$.

If we apply the limit functor Lim to this diagram D (of diagrams) with base I, and then take its limit, we get the following diagram in $\mathscr{T}$,

$$L \longrightarrow \Pi_j\ \mathscr{O}(U_j) \rightrightarrows \Pi_{j,j'}\ \mathscr{O}(U_j \cap U_{j'}).$$

The equaliser formulation of the sheaf condition (Definition 2.20) says that $\mathscr{O}$ is a sheaf iff $L \cong \mathscr{O}(U)$ whenever $U = \bigcup_j U_j$, and the Behaviour Theorem says that L is the limit of the following diagram



where all arrows are restriction morphisms. From this, we conclude that $\mathcal{O}$ is a sheaf iff its behaviour $\mathcal{O}(U)$ at U is the limit of its behaviours $\mathcal{O}(U_j)$ at U_j, whenever $U = \bigcup_j U_j$. For example, if $\mathcal{T} = \mathcal{I}_0(\omega)$ and $U_j = [n]$, then we have $\omega = \bigcup_n [n]$, so that $\mathcal{O}(\omega)$ is the limit of the $\mathcal{O}([n])$.

3.4. *Discussion*

We have considered systems at three different levels: (1) *objects*, as collections of possible observations; (2) *systems*, as collections of interacting objects; and (3) *interconnections*, as systems of interacting systems. A wide variety of systems can be treated in this way, including digital hardware, electrical circuits, and (as shown in the next section) concurrent programming languages. Concepts from category theory have helped achieve this generality: we model objects as sheaves; systems as diagrams; and interconnections as diagrams of diagrams. In addition, behaviour is given by limit, and the result of interconnection by colimit. Although we have not done so here, it is possible to iterate these constructions to obtain hierarchical systems of arbitrary depth; see Goguen (1971, 1973), Goguen and Ginali (1978).

Our approach to systems is *declarative* or *constraint based* in the sense that behaviour arises through "mutual effects" or "interdependent origination" rather than through the propagation of causes and effects; in particular, we do not assume that all devices have inputs and outputs, and hence we are not limited to simple functional devices[7]. In this setting, what it means for a system to *satisfy* a specification (i.e., a presheaf representing some property, such as fairness) is that if we interconnect that property with the system, then the resulting behaviour is the same as (i.e., is isomorphic to) the behaviour of the original system.

It is interesting to look at so-called "internal choice" and "external choice" in the context of limits of diagram of sheaves. We have already noted that non-determinism simply corresponds to there being more than one function to choose from in some set $\mathcal{O}_n(U)$. When $\mathcal{O}_n$ participates in some larger system, some of these elements may no longer be consistent with overall system behaviour – i.e., they may not be in the n component of any consistent net of points for the system. More interestingly, the n component may be completely determined by the behaviour of the rest of the system. In this case, we may say that "external choice" is being exercised. But if several values for the nth component remain, we may say that $\mathcal{O}_n$ has "internal choice". Under a constraint oriented view of systems, the distinction between internal and external choice appears somewhat artificial, and may depend on the point of view taken. (Formulating the above discussion for an arbitrary structure category requires some additional concepts that are discussed in Section 4.2.)

It is worth noting that the sheaves that arise from some particular model of concurrency, such as some kind of transition system, typically form a complete subcategory of the

[7] Of course, ours is not the only approach with this property. For example, Gordon (1986) has used higher order relations to study digital circuits. But the point does seem worth a bit of emphasis.

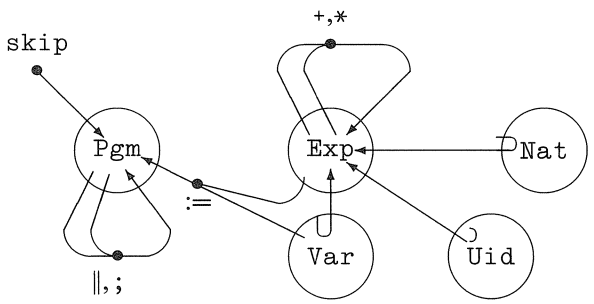
category of all sheaves over the appropriate base; in general, there will be some sheaves that do not correspond to any system of the given kind.

4. Semantics and Properties of Systems

Now that we have stated, motivated and illustrated all four of the main principles of this paper, let us apply them to some examples of more immediate interest to Computing Science. This section shows how to give semantics for a simple concurrent language executing in a distributed environment.

The objects considered in the previous sections were *closed*, in the sense that their attributes can be divided into inputs and outputs such that the values on the inputs uniquely determine those on the outputs. But modern programming requires *open systems*, the objects of which only constrain what happens under certain conditions, and leave the rest unconstrained. (Of course, non-determinism, where outputs are only partially constrained, is also possible.) Openness in this sense is similar to what happens when a number of objects are connected to an Ethernet, and each responds only to those messages addressed to it.

The syntax of our simple concurrent programming language is given by the (order-sorted) signature[8] Σ of the following ADJ diagram:



Here `Uid` is the sort for object names, `Var` for variables, `Nat` for natural numbers, `Exp` for expressions, and `Pgm` for programs. Also, + and * are operations on expressions, while ; and || are operations on programs, := is assignment, `skip` is the instruction that does nothing, and $\hookrightarrow$ indicates a *subsort* relation (based on order-sorted algebra (Goguen and Meseguer 1989)). We also assume that the sorts `Uid`, `Var` and `Nat` are populated by a countably infinite number of mutually disjoint constants. Typical elements of sort `Uid` are `Tom`, `Dick` and `Harry`, and of `Var` are X, Y, Z, while the elements of sort `Nat` are of course 0,1,2,... . We let $|\Sigma|$ denote the set of all symbols in Σ; hence $|\Sigma|^*$ denotes the set of all finite strings of elements from Σ. Unless otherwise indicated, variables w, w', w_1, etc. range over $|\Sigma|^*$.

[8] An order-sorted signature has, in addition to the data of a many sorted signature, a partial ordering on the set S of sorts, and some assumptions about the consistency of overloaded operation symbols. This ADJ diagram is therefore augmented to indicate subsort relationships with "hooked" arrows. (See Goguen and Meseguer 1989) for details of order-sorted algebra.)

Our language is the algebra T_Σ of all Σ-terms, constructed just like the order-sorted term algebra in Goguen and Meseguer (1989), except that in order to simplify the transition rules to be given below, we will use reverse Polish (i.e., Polish postfix, or Łukasiewicz) notation for the terms of sort `Pgm`, which will represent programs. For example, the program

$$X := Y + 2 \parallel Y := X + 2$$

appears in $T_{\Sigma,\mathtt{Pgm}}$ as the string

$$\parallel := X\ Y\ 2\ +\ := Y\ X\ 2\ +$$

To simplify the notation, we will sometimes write $u \in \mathtt{Uid}$ instead of $u \in T_{\Sigma,\mathtt{Uid}}$ and similarly for $X \in \mathtt{Var}$ and the other sorts. So $n \in \mathtt{Nat}$ iff $n \in \omega$. Also, we abbreviate $T_{\Sigma,\mathtt{Pgm}}$ by P_Σ.

For simplicity, we will use discrete time in this example, so that $\mathscr{T} = \mathscr{I}_0(\omega)$. Then the appropriate event stream object $\mathscr{E}$ is given, for $I \in \mathscr{T}$, by

$$\mathscr{E}(I) = \{f \mid f : I \to P_\Sigma\}.$$

Next, we give the objects that define the semantics of the various features of the language. For this purpose, we will use the following "transition notation" in defining objects $\mathscr{O}$ over the base $\mathscr{I}_0(\omega)$: if A is the attribute object of $\mathscr{O}$, then $a \mapsto a'$ (for $a, a' \in A$) means that if $f(i) = a$ for some $f \in \mathscr{O}(I)$, then there are some $I' \supseteq I$ and $f' \in \mathscr{O}(I')$ such that $i+1 \in I'$ and $f' \restriction I = f$ and $f'(i+1) = a'$. (This implies that each basic operation takes unit time; but it is easy to redefine $\mapsto$ so that different operations take different times, or so that operations may take as long as they like.) In general, a and a' in the notation $a \mapsto a'$ are not elements of A, but rather may be *patterns* that define a set of transitions. For example,

$$w{+}nmw' \mapsto wkw' \quad \textbf{if} \quad n, m \in \omega \quad \textbf{and} \quad k = n + m$$

means that if there exist strings $w, w' \in |\Sigma|^*$ and some numbers n, m such that $f(i) = w{+}nmw'$, then there is some f' such that $f'(i) = w{+}nmw'$ and $f'(i+1) = wkw'$ where $k = n + m$. Also, we will let

$$\mapsto a$$

indicate that the initial state is a, i.e., that if $f \in \mathscr{O}(I)$ is defined at 0, then $f(0)$ matches the pattern a; usually a is just a single attribute in this notation.

The object $\mathscr{O}$ defined by a set of transition relations is then the least family[9] $\mathscr{O}(U) \subseteq \mathscr{E}(U)$ of sets of term-valued functions, for each $U \in \mathscr{T}$, satisfying the transition relations and closed under restriction, i.e., such that if $U \subseteq V$ and $f \in \mathscr{O}(V)$ then $f \restriction U \in \mathscr{O}(U)$.

We now define a series of objects, one for each feature of the language, which when put together give a system whose semantics is that of the language.

[9] This exists because the conditions defining it are all positive, i.e., they are Horn clauses.

Example 4.1. (Parallel Composition) This feature is defined by an object denoted $\mathscr{P}$ having attribute object P_Σ and satisfying

$$\begin{array}{lll} w_1 \parallel \texttt{skip}\ w\ w_2 & \mapsto & w_1'\ w\ w_2' \\ w_1 \parallel w\ \texttt{skip}\ w_2 & \mapsto & w_1'\ w\ w_2' \end{array}$$

for $w_1, w_1', w, w', w_2, w_2' \in P_\Sigma$. This just says that when either of the two parallel programs is completed, the parallel constructor can be eliminated; computation by the other program can proceed in parallel with the elimination. Note that this object is a subobject of $\mathscr{E}$; we will use the resulting inclusion morphism to relate it to other objects.

Example 4.2. (Sequential Composition) This object is denoted $\mathscr{Q}$, has attribute object P_Σ again, and is defined by

$$\begin{array}{lll} w_1\ ;w_2\ w_3\ w_4 & \mapsto & w_1'\ ;w_2'\ w_3\ w_4' \ \textbf{if}\ w_1 \neq \texttt{skip} \\ w_1\ ;\texttt{skip}\ w_3\ w_4 & \mapsto & w_1'\ w_3\ w_4' \end{array}$$

This says that the first program must be completed before the execution of the second can begin. It does not constrain what can be done by the first program. Again, its projection is the inclusion.

Example 4.3. (Assign) Assignment is defined by a family of objects $\mathscr{G}^X$, one for each $X \in \texttt{Var}$. $\mathscr{G}^X$ has attribute object $\omega \times P_\Sigma$ and is defined by the following:

$$\begin{array}{lll} & \mapsto & \langle 0, w\rangle \\ \langle m, w := Xnw'\rangle & \mapsto & \langle n, w''\texttt{skip}\,w'''\rangle \\ \langle n, wpXw'\rangle & \mapsto & \langle n, w''pnw'''\rangle \ \textbf{if}\ p \neq := \end{array}$$

where $n \in \omega$ and $p \in \Sigma$. This says the each variable initially has value 0, that an assignment of n to X causes the object $\mathscr{G}^X$ to remember n, and that the variable X can be replaced by its value, unless it occurs just after $:=$. The projection is from the second component of the pair in the state.

Example 4.4. (Adder) Let $\omega_\perp = \omega \cup \{\perp\}$ and let $\omega_\perp \times P_\Sigma$ be the attribute object. Then a typical adder object, $\mathscr{A}^{\texttt{Fred}}$ for $\texttt{Fred} \in \texttt{Uid}$, may be defined as follows:

$$\begin{array}{lll} \langle s, w_1\texttt{Fred}\,w_2\rangle & \mapsto & \langle \perp, w_1' s w_2'\rangle \ \textbf{if}\ s \neq \perp \\ \langle \perp, w_1 {+} nmw_2\rangle & \mapsto & \langle s, w_1'\texttt{Fred}\,w_2'\rangle \ \textbf{where}\ s \text{ is the number } n{+}m \\ \langle s, w\rangle & \mapsto & \langle s, w'\rangle \ \textbf{if}\ w, w' \in P_{\Sigma\text{–}\texttt{Fred}} \end{array}$$

where $n, m \in \omega$. This says that Fred performs just one addition at a time, and allows other operations to occur concurrently. If additions are ready to be performed, then he must perform one; he replaces the chosen expression $+nm$ by his name and memorises the sum. When he sees his name, he replaces it by the number he has memorised, forgets that number, and becomes ready to do another sum. Notice that Fred also prohibits anyone else (to whom he communicates) from using his name. Again, projection is from the second component of the pair in the state.

Example 4.5. (Multiplier) A typical multiplier, say $\mathscr{M}^{\texttt{Tom}}$ for $\texttt{Tom} \in \texttt{Uid}$, is defined just like the adder Fred, except that + is replaced by $*$ and Fred is replaced by Tom.

We have now defined all the features of the language. Unfortunately, what we have still

allows unprogrammed transitions to occur. However, we can introduce a "frame object" to keep things the same unless they are deliberately changed:

Example 4.6. (Frame) This object is denoted $\mathscr{F}$, has attribute object P_Σ, and is defined by

$$w_1 ... w_n \mapsto w_1' ... w_n' \quad \textbf{where}$$
$$w_i = w_i' \text{ for } i = 1, ..., n \quad \textbf{unless}$$
$$w_i \in \texttt{Uid} \quad \textbf{or} \quad w_i' \in \texttt{Uid} \quad \textbf{or} \quad w_i \in \{:=, \|, ;, \texttt{skip}\},$$

where w_i, w_i' are here in Σ^* rather than in P_Σ.

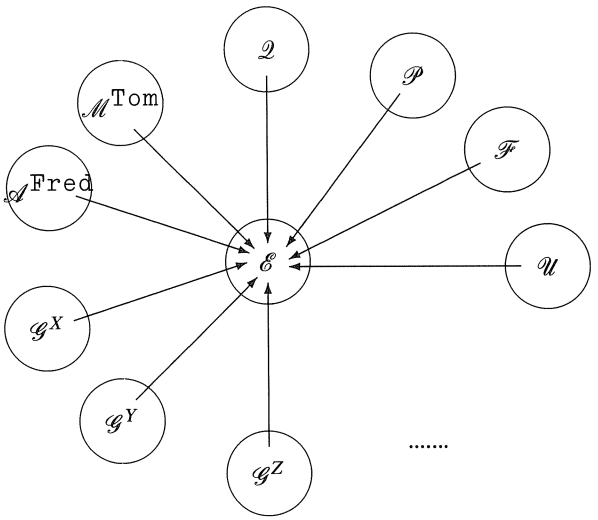
Finally, in order to "close" the system, we can add a "shop steward" who disallows unregistered workers:

Example 4.7. (Union) This object, denoted $\mathscr{U}$, is defined, for $I \in \mathscr{T}$, by

$$\mathscr{U}(I) = \{f \mid f : I \to P_{\Sigma - (\texttt{Uid} - W)}\},$$

where W is the set of union members, which might for example be {Fred, Tom}.

So our system looks as follows,



and the limit of this diagram evaluates programs in our little language. In the limit object, at each instant of time, each object has a copy of the same program, and possibly some internal state, such as the value of a variable. As time progresses, the program is simplified as worker objects process its parts. This processing is concurrent and distributed. Moreover, the objects for variables are true objects in the sense of object oriented programming, although very simple ones. The consistent nets of points in the limit object can be seen as the *run time states* of the computation. If we add more workers, programs can be executed more quickly. Note that these objects can be seen as (infinite state) automata.

Of course, this is a simple example but it seems clear that the same techniques will extend to much more complex languages. For example, it is easy to add more language constructs, such as loops. Another interesting feature to add would be `abort`. Also, it

seems that the semantics of the functional and object oriented language FOOPS (Goguen and Meseguer 1987, Goguen and Wolfram 1991) and of the Rewrite Rule Machine (see Goguen *et al.* 1989 and Goguen 1991b) can be developed in a similar way, and I hope these will be discussed in future papers. See Ehrich *et al.* (1991) for some other applications of sheaf theory to concurrent systems.

4.1. *Deadlock*

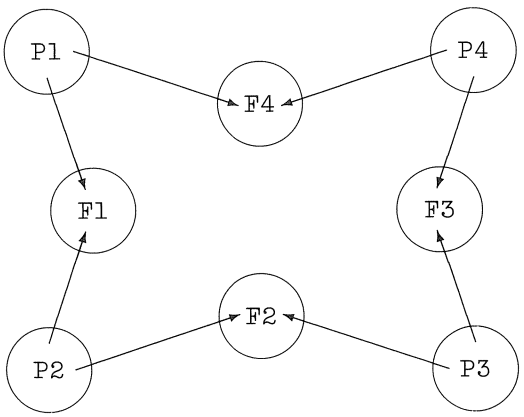
It is important that many real systems should be *deadlock free*, in the sense that they do not get into dead states. Definitions of deadlock in the literature tend to be syntactic, and as far as I am aware, there is no definition that is sufficiently general to encompass all of the kinds of system considered in this paper. The following proposes one:

Definition 4.8. If $\mathscr{L}$ is the limit object of a system, then the system has **deadlock** at $h \in \mathscr{L}(U)$ iff for all $U' \supseteq U$ and all $h' \in \mathscr{L}(U')$ such that $h' \restriction U = h$, we have $h'(i) = h'(i')$ for all $i, i' \in U' - U$. Furthermore, the system **terminates** at $h \in \mathscr{L}(U)$ iff $(\forall U' \supseteq U)(\forall h' \in \mathscr{L}(U'))\ h' \restriction U \neq h$.

Deadlock at h says that if the system has evolved through the events in h, then evolution into a different state is impossible. Terminating at h is a more drastic form of deadlock in which there is no possible future behaviour of the system that extends h. For example, consider a short circuited 6 volt battery: it cannot have both a 6 and a 0 volt potential difference between its terminals. Thus, a discrete time system in a state where in the next instant a switch will be thrown that produces a short circuit is terminal at that state. (Of course, a more realistic model would show a large current flow followed by the decay of either the battery and/or the wiring, but in the simple model chosen for this example, a short circuit is an inconsistency which precludes any future behaviour.)

The following is a common example in the Computing Science literature on concurrency:

Example 4.9. (Dining Philosophers): This somewhat fanciful situation involves four philosophers supported in a research institute with a circular table, the centre of which always contains a plate of food. This food must be eaten seated at the table with one fork in each hand. The table has four forks, one between each two adjacent chairs. Philosophers are asynchronous processes that think and eat.
Let us now construct a formal model of this situation: let $\mathscr{P} = \{P1, P2, P3, P4\}$ be the names of the four philosophers, and let $\mathscr{F} = \{F1, F2, F3, F4\}$ be names for the four forks. Then the following is a diagram for this situation as a system:



The fork object between philosophers P and P' has states $s \in \{P, P', \bot\}$ with the following transitions,

$$\begin{array}{lll} s & \mapsto & s \\ \bot & \mapsto & s \\ s & \mapsto & \bot \end{array}$$

That is, a fork can remain as it is, can be picked up, and can be put down.
Each philosopher P has states $s = \langle l, n, r\rangle$ for $l, r \in \mathscr{P}_\bot = \mathscr{P} \cup \{\bot\}$ and $n \in \{t, e\}$ (t and e stand for "thinking" and "eating"[10] respectively), and has the following transitions:

$$\begin{array}{lcll} & \mapsto & \langle \bot, t, \bot\rangle & \\ \langle l, n, r\rangle & \mapsto & \langle l', n, r'\rangle & \textbf{if } l, l', r, r' \neq P \\ \langle l, t, r\rangle & \mapsto & \langle l', e, r'\rangle & \textbf{if } l, l', r, r' \neq P \\ \langle l, e, \bot\rangle & \mapsto & \langle l', e, P\rangle & \textbf{if } l, l' \neq P \\ \langle l, e, P\rangle & \mapsto & \langle l', e, P\rangle & \textbf{if } l, l' \neq P \\ \langle \bot, e, P\rangle & \mapsto & \langle l, e, P\rangle & \\ \langle P, e, P\rangle & \mapsto & \langle \bot, t, \bot\rangle & \end{array}$$

The first rule says that philosophers are born thinking, without forks. The second rule says that if a philosopher has no forks, then he can continue in his present state, without constraining the actions of his two neighbouring philosophers. The third rule says if he has no forks and is thinking, then he may become hungry, again without constraining his neighbours. The fourth rule says that if he has no forks, then he can pick up his right fork if it is available, without constraining his left neighbour (note that, by the second rule, he need not do so). The fifth rule says that he can remain in the state of having just a right fork, without constraining his left neighbour (but he does not need to so remain, because l can be any fork state). The sixth rule says that if he has a right fork, then he can pick up his left fork, if it is available. The last rule says that if he has both forks and is eating, then he can put them both down and think. (A slightly more accurate model

[10] A more accurate description would be "hungry and in a process that may lead to eating."

might add a rule saying that if he has both forks he can continue to hold them and eat; under the above rules, he has only one unit of time in which to eat.)
The object for each philosopher has two projections, from its first and third components, as shown in the diagram above.
Let $\mathscr{L}$ denote the limit of this system, and let $I = \{0, ..., i\} \in \mathscr{I}_0(\omega)$. Then officially, the elements of $\mathscr{L}(I)$ are 8-tuples of functions on I, four of them 3-tuples for the philosophers, and four of them elements of $\mathscr{P}_\perp$, for the forks. However, it is equivalent to consider 16-tuple valued functions on I, by flattening the 3-tuples. In fact, these 16-tuples have a lot of redundancy, due to the constraints imposed by their representing consistent nets of points, and it suffices to consider 8-tuples of the form

$$s = \langle f_1, f_2, f_3, f_4, p_1, p_2, p_3, p_4 \rangle$$

with $f_i \in \mathscr{P}_\perp$ and $p_i \in \{t, e\}$, where f_i is the state of the i^{th} fork and p_i is the middle component of the state of the i^{th} philosopher. Hereafter, we will feel free to call p_i "the state of P_i" and to call such 8-tuples "states of the system". Notice that when the system is in state s, each philosopher P_i (for $i = 1, 2, 3, 4$) is in the state $\langle f_{i-1}, p_i, f_{i+1} \rangle$, where the subscripts are understood modulo 4.
There is a unique state of the system in which each philosopher is in his initial state, namely

$$\langle \perp, \perp, \perp, \perp, t, t, t, t \rangle$$

and we will call this the "initial state of the system".
It is now easy to see that the state in which each philosopher holds his right fork is a reachable deadlock state. To see that it is reachable, it suffices to give a sequence of transitions from the initial state to it. In fact, two parallel transitions are enough: in the first, each philosopher becomes hungry, and in the second, each philosopher picks up his right fork:

$$\langle \perp, \perp, \perp, \perp, t, t, t, t \rangle \mapsto \langle \perp, \perp, \perp, \perp, e, e, e, e \rangle \mapsto \langle P_1, P_2, P_3, P_4, e, e, e, e \rangle .$$

Let us denote this last state by x. To see that it is a deadlock state, we note that only the second rule applies to x. This is checked by seeing which lefthand sides of rules match $\langle P_{i-1}, e, P_{i+1} \rangle$, where the subscripts are again understood modulo 4. Therefore, the only possible transitions from x return the system to state x, because the corresponding righthand side of the second rule is again $\langle P_{i-1}, e, P_{i+1} \rangle$. This analysis also shows that the system is not terminal.

It is possible to modify this example to use real time, by giving a slightly different interpretation for the transition notation $\mapsto$; however, we omit this.

Deadlock can also arise when testing to see if some system satisfies some property by interconnecting the system with the property and taking the limit: if the property is *inconsistent* with the system, in the sense that the system has no states with the given property, then the interconnected system exhibits deadlock.

4.2. *Information Flow and Security*

Another important property of real systems is *security*: we may want to be sure that access to certain confidential information is protected, and that certain unauthorised actions are prevented. For example, we don't want an electronic bank robber to be able to discover which are the largest accounts, and then withdraw funds from them. It is now rather well recognised that "non-interference" assertions can be used to express security properties of (monolithic) sequential systems (Goguen and Meseguer 1982, 1984). An active research topic is the extension of such assertions to more general classes of system. Definition 4.12 below extends non-interference to a much wider class of system, such as distributed object oriented databases, and appears to be more general than anything else in the literature. (But see Jacob (1988), O'Halloran (1990) for two other very general approaches.) First, we need some auxiliary material:

Definition 4.10. Let $\varphi : \mathcal{O} \to \mathcal{O}'$ be a morphism of set valued presheaves. Then the **image presheaf** of φ, denoted $\varphi(\mathcal{O}) \subseteq \mathcal{O}'$, is defined by $\varphi(\mathcal{O})(U) = \varphi_U(\mathcal{O}(U))$; note that $\varphi(\mathcal{O})(U) \subseteq \mathcal{O}'(U)$ for each $U \in \mathcal{T}$.

This definition can be generalised to a suitable structure category $\mathcal{C}$, and is used in the last equation of Definition 4.12 below. But first note the following:

Fact 4.11. The image of a morphism of sheaves is a sheaf.

Now the main concept:

Definition 4.12. Given a system $\mathcal{S}$ with objects $\mathcal{S}_n$ for $n \in N$, with morphisms $\varphi_e : \mathcal{S}_n \to \mathcal{S}_{n'}$ for $e : n \to n'$, and with behaviour (i.e., limit) $\mathcal{L}$ having projections $\pi_n : \mathcal{L} \to \mathcal{S}_n$, then $\mathcal{S}_m$ is **non-interfering with** $\mathcal{S}_k$, written $\mathcal{S}_m \not\rightsquigarrow \mathcal{S}_k$, iff the following holds:

let $\mathcal{L}'$ be the limit of the subsystem of $\mathcal{S}$ from which $\mathcal{S}_m$ and all morphisms to and from $\mathcal{S}_m$ have been omitted; let $\pi'_n : \mathcal{L}' \to \mathcal{S}_n$ be its projections, for $n \in N - \{m\}$; then $\pi_k(\mathcal{L}) = \pi'_k(\mathcal{L}')$.

What this says[11] is that the behaviour of the system looks the same from the object $\mathcal{S}_k$ with the object $\mathcal{S}_m$ omitted as it does with the object $\mathcal{S}_m$ present; i.e., there is no flow of information from $\mathcal{S}_m$ to $\mathcal{S}_k$. This definition is general enough to apply to data dependency analysis for compilers of concurrent languages onto distributed systems, and to the flow of information in natural language conversation, along the lines suggested by situation semantics (Barwise and Perry 1983) and the work of Dretske (1981). Indeed, it seems possible that sheaf theory could help in providing a natural semantics for situation theory.

5. Conclusions

Our sheaf approach provides a semantic, i.e., model theoretic, foundation for concurrent distributed computing by (possibly active) objects, without commitment to any particular notation or conceptualisation for concurrency. In general, such an approach should be closer to our physical intuition, and can provide standards against which to measure the soundness and completeness of syntactic systems.

[11] It may help to think of "k" as a "crook" and "m" as a "market analyst" whose information k seeks to steal.

Sheaf theory has been used in mathematics to study relationships between local and global phenomena, for example, in algebraic geometry, differential geometry, and even logic; the subject has also been developed in an abstract form using category theory. The theory of topoi, originally developed by Lawvere and Tierney (see Lawvere 1975, Goldblatt 1979, Barr and Wells 1984) is perhaps the most exciting development in this respect. An interesting topic for future research is to see what the theory of topoi can tell us about concurrency. For example, one should be able to reason about a system using the internal intuitionistic logic of the corresponding topos of sheaves.

Concepts from category theory have helped us achieve generality: objects have been modelled by sheaves; inheritance by sheaf morphisms; systems by diagrams; and interconnections by diagrams of diagrams. In addition, behaviour is given by limit, and the result of interconnection by colimit. The approach is illustrated with many examples, including a semantics for a simple concurrent object-based programming language.

The definitions, examples and results in this paper are just a beginning. Yet the variety of examples may be surprising. Moreover, one example illustrates an important class of applications, namely the semantics of concurrent, distributed object oriented systems. Some of the definitions may also be surprising for their generality, including a notion of security that generalises the Goguen-Meseguer non-interference approach (Goguen and Meseguer 1982, 1984) from sequential systems to (for example) real time distributed concurrent object oriented databases. A very general definition of deadlock is also given. It is interesting that these concepts are so easily stated in a purely semantic form.

Acknowledgements

Special thanks to Rod Burstall for "active listening" at the Dôme Café in Oxford while I tried to explain these ideas in preparation for a lecture on this subject at the UK-Japan Workshop on Concurrency, in September 1989. Thanks also to Hans-Dieter Ehrich, Amilcar Sernadas, José Fiadeiro, Felix Costa, Tom Maibaum and others in the ESPRIT sponsored ISCORE project (e.g., see Ehrich and Sernadas 1990, Ehrich *et al.* 1989, and Ehrich *et al.* 1990) for reawakening my interest in this area, and for their encouraging comments. Thanks to Susanna Schwab (neé Ginali) for a number of very useful suggestions, thanks to Jeremy Jacob and David Wolfram for noticing a number of bugs and infelicities in an earlier draft, and thanks to David Benson for showing me drafts of a paper in preparation, and the MSc thesis of Rakesh Dubey, both of which influenced the final draft of this paper. Thanks to Răzvan Diaconescu for several valuable suggestions, particularly the proof of Theorem 3.9. Thanks also to Frances Page and Joan Arnold for help with preparation of the manuscript, and especially with the diagrams.

This paper is dedicated with affection to Professor Erwin Engeler, and was given as a lecture on the occasion of his sixtieth birthday, in March 1990.

REFERENCES

Agha, G. (1986). *Actors: A Model of Concurrent Computation in Distributed Systems.* MIT.

Barr, M. and Wells, C. (1984). *Toposes, Triples and Theories.* Springer. Grundlehren der mathematischen Wissenschafter, Volume 278.

Barr, M. and Wells, C. (1990). *Category Theory for Computing Science.* Prentice-Hall.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes.* MIT. Bradford Books.

Bergstra, J. and Klop, J.W. (1984). Process algebra for synchronous communication. *Information and Control*, 60:190–137.

Dretske, F. (1981). *Knowledge and the Flow of Information.* MIT. Bradford Books.

Ehrich, H.D., Goguen, J. and Sernadas, A. (1991). A categorial theory of objects as observed processes. In J.W. deBakker, Willem P. de Roever, and Gregorz Rozenberg, editors, *Foundations of Object Oriented Languages*, pages 203–228. Springer. Lecture Notes in Computer Science, Volume 489; Proceedings, REX/FOOL Workshop, Noordwijkerhout, The Netherlands, May/June 1990.

Ehrich, H.D. and Sernadas, A. (1990). Algebraic implementation of objects over objects. In J.W. deBakker, Jan Willem deRoever, and Gregorz Rozenberg, editors, *Proceedings, REX Workshop on Stepwise Refinement of Distributed Systems: Models, Formalisms, Correctness*, pages 239–266. Springer. Lecture Notes in Computer Science, Volume 430.

Ehrich, H.D., Sernadas, A. and Sernadas, C. (1989). Objects, object types, and object identification. In Hartmut Ehrig et al., editors, *Categorical Methods in Computer Science with Aspects from Topology*, pages 142–156. Springer. Lecture Notes in Computer Science, Volume 393.

Ehrich, H.D., Sernadas, A. and Sernadas, C. (1990). From data types to object types. *Journal of Information Processing and Cyberntics*, 26(1/2):33–48.

Ferrari, G.L. (1990). *Unifying Models of Concurrency.* PhD thesis, University of Pisa.

Girard, J.Y. (1986). Linear logic. *Theoretical Computer Science*, 46:1–102.

Goguen, J. (1971). Mathematical representation of hierarchically organized systems. In E. Attinger, editor, *Global Systems Dynamics*, pages 112–128. S. Karger.

Goguen, J. (1973). Categorical foundations for general systems theory. In F. Pichler and R. Trappl, editors, *Advances in Cybernetics and Systems Research*, pages 121–130. Transcripta Books.

Goguen, J. (1975). Objects. *International Journal of General Systems*, 1(4):237–243.

Goguen, J. (1991a). A categorical manifesto. *Mathematical Structures in Computer Science*, 1(1):49–67. Also, Programming Research Group Technical Monograph PRG-72, University of Oxford, March 1989.

Goguen, J. (1991b). Semantic specifications for the rewrite rule machine. In Aki Yonezawa and Takayasu Ito, editors, *Concurrency: Theory, Language and Architecture*, pages 216–234. *Proceedings of a U.K.–Japan Workshop.* Springer. Lecture Notes in Computer Science, Volume 491.

Goguen, J. and Ginali, S. (1978). A categorical approach to general systems theory. In George Klir, editor, *Applied General Systems Research*, pages 257–270. Plenum.

Goguen, J., Leinwand, S., Meseguer, J., and Winkler, T. (1989). *The Rewrite Rule Machine, 1988.* Technical Report Technical Monograph PRG-76, Programming Research Group, Oxford University.

Goguen, J. and Meseguer, J. (1982). Security policies and security models. In Marvin Schafer and Dorothy D. Denning, editors, *Proceedings, 1982 Symposium on Security and Privacy*, pages 11–22. IEEE Computer Society.

Goguen, J. and Meseguer, J. (1984). Unwinding and inference control. In Dorothy D. Denning and Jonathan K. Millen, editors, *Proceedings, 1984 Symposium on Security and Privacy*, pages 75–86. IEEE Computer Society.

Goguen, J. and Meseguer, J. (1987). Unifying functional, object-oriented and relational programming, with logical semantics. In Bruce Shriver and Peter Wegner, editors, *Research Directions in Object-Oriented Programming*, pages 417–477. MIT. Preliminary version in *SIGPLAN Notices*, Volume 21, Number 10, pages 153-162, October 1986.

Goguen, J. and Meseguer, J. (1989). Order-sorted algebra I: Equational deduction for multiple inheritance, overloading, exceptions and partial operations. Technical Report SRI-CSL-89-10, SRI International, Computer Science Lab. Given as lecture at Seminar on Types, Carnegie-Mellon University, June 1983; many draft versions exist.

Goguen, J., Thatcher, J. and Wagner, E. (1976). An initial algebra approach to the specification, correctness and implementation of abstract data types. Technical Report RC 6487, IBM T.J. Watson Research Center. Also in *Current Trends in Programming Methodology, IV*, Raymond Yeh, editor, Prentice-Hall, 1978, pages 80-149.

Goguen, J. and Wolfram, D. (1991). On types and FOOPS. In Robert Meersman and Samjit Khosla, editors, *Proceedings, IFIP TC2 Conference on Object Oriented Databases.* North Holland, to appear.

Goldblatt, R. (1979). *Topoi, the Categorial Analysis of Logic.* North-Holland.

Gordon, M.J.C. (1986). Why higher-order logic is a good formalism for specifying and verifying hardware. In George Milne and P.A. Subrahmanyam, editors, *Formal Aspects of VLSI Design.* North-Holland.

Gray, J. (1965). Sheaves with values in a category. *Topology*, 3(1):1–18.

Grothendieck, A. (1963). Catégories fibrées et descente. In *Revêtements étales et groupe fondamental, Séminaire de Géométrie Algébraique du Bois-Marie 1960/61, Exposé VI.* Institut des Hautes Études Scientifiques. Reprinted in Lecture Notes in Mathematics, Volume 224, Springer, 1971, pages 145–94.

Hoare, C.A.R. (1985). *Communicating Sequential Processes.* Prentice-Hall.

Jacob, J. (1988). A security framework. In *Proceedings, 1988 Computer Security Foundations Workshop*, pages 98–111. MITRE. Franconia, New Hampshire.

Mac Lane, S. (1971). *Categories for the Working Mathematician.* Springer.

Lawvere, F.W. (1975). Introduction. In F. William Lawvere, C. Maurer, and Gavin Wraith, editors, *Model Theory and Topoi*, pages 3–14. Springer. Lecture Notes in Mathematics, Volume 445.

Lilius, J. (1991). *Sheaf semantics for Petri nets.* Technical report, Helsinki University of Technology.

Meseguer, J. and Montanari, U. (1988). Petri nets are monoids: A new algebraic foundation for net theory. In *Proceedings, Symposium on Logic in Computer Science.* IEEE Computer Society. Full version in Technical Report SRI-CSL-88-3, Computer Science Laboratory, SRI International, January 1988; submitted to *Information and Computation.*

Milner, R. (1980). *A Calculus of Communicating Systems.* Springer. Lecture Notes in Computer Science, Volume 92.

Monteiro, L. and Pereira, F. (1986). *A sheaf-theoretic model of concurrency.* Technical Report CSLI-86-62, Center for the Study of Language and Information, Stanford University.

O'Halloran, C. (1990). *A calculus of information flow.* Technical report, Royal Signals and Radar Establishment, Malvern.

Pierce, B.C. (1990). *A taste of category theory for computer scientists.* Technical Report CMU-CS-90-113, Carnegie-Mellon University.

Reisig, W. (1986). *Petri Nets: An Introduction.* Springer. EATCS Monographs on Theoretical Computer Science.

Star, S.L. (1988). The structure of ill-structured solutions: heterogeneous problem-solving, boundary objects and distributed artificial intelligence. In Michael Huhns and Les Gasser, editors, *Distributed Artificial Intelligence*, volume 3, pages 37–54. Morgan Kauffmann.

Stavridou, V., Goguen, J., Eker, S. and Aloneftis, A. (1991). FUNNEL: A CHDL with formal semantics. In *Proceedings, Advanced Research Workshop on Correct Hardware Design Methodologies*, pages 117–144. IEEE. Turin.

Tarlecki, A., Burstall, R., and Goguen, J. (1991). Some fundamental algebraic tools for the semantics of computation, part 3: Indexed categories. *Theoretical Computer Science*, 91:239–264. Also, Technical Report PRG-77, August 1989, Programming Research Group, Oxford University.

Tennison, B.R. (1975). *Sheaf Theory*. Cambridge. London Mathematical Society Lecture Notes Series, 20.

Winskel, G. (1990). A compositional proof system on a category of labelled transition systems. *Information and Computation*, 87:2–57.

Yosida, K. (1968). *Functional Analysis*. Springer. Second Edition.