The Tblisi Symposium on Logic, Language and Computation

edited by Jonathan Ginzburg et al.

Contents

v

1

A Generalizable Semantics for a Default Inheritance Reasoner

Carl Vogel

ABSTRACT. Path-based default inheritance reasoning can be seen as a subset of the monadic predicate calculus with a guarantee of inferential acyclicity and a default interpretation of conditionals. This paper offers a channel theoretic semantics for the theory of skeptical inheritance developed by Horty et al. 1990. It takes analysis of constraints given by channel theory, one which makes a distinction between type and token aspects of both constraints and the situations which are related by them. The particular formulation of channel composition given here provides an interpretation to paths prefered by the skeptical inheritance reasoner mentioned. The semantics is generalizable in that interpretations of related reasoners can be obtained simply by varying the formulation of channel composition.

1 Introduction

Inheritance networks provide a formal tool for reasoning about classifications of individuals and concepts as an alternative to first order logic (FOL), providing a form of paraconsistency in which direct contradictions can be contained in a theory without rendering it useless. Relationships between individuals and classes are stated in terms of directed acyclic graphs of *is-a* and *is-not-a* relations, but other relations are possible (Brachman 1983, Touretzky 1986). The relations may be strict or defeasible and networks of links that represent those relations may contain both sorts. By reasoning over such a network (via constructing paths through it), one may determine the properties possessed by an individual. For a given network, there are numerous reasoning methods which differ in the manner in which they treat conflicting information (i.e. which paths are preferred). The differences are classified into a number of axes of variation by Touretzky et al.

The Tblisi Symposium on Logic, Language and Computation
J. Ginzburg et al., eds.

1987 and Vogel 1995. While Touretzky 1986 provided a lattice based semantics for his first system, it has been pointed out that a great deal of the literature focuses on proof theory rather at the expense of general semantics for inheritance reasoners (e.g. Boutilier 1989). In fact, semantics has been provided by a number of researchers for individual reasoning systems (Delgrande 1990, Boutilier 1989), but no formal semantics has been outlined which provides a general enough framework to capture a large class of reasoning systems. This paper presents a semantics for the inheritance reasoner presented by Horty et al. 1990 (hereafter, H90). This semantics is quite general; Vogel 1995 demonstrates how the to parameterize the semantics to construct models for a family of path-based inheritance reasoners defined by Touretzky 1986, Touretzky et al. 1987, Horty and Thomason 1988. The semantics uses tools for talking about natural constraints provided by situation theory and channel theory.

2 Modeling Tools

The theory of constraints from situation theory (Barwise 1993) augmented with the theory of information channels from channel theory (Seligman 1990, Seligman and Barwise 1993, 1994, Barwise 1993, Barwise and Seligman 1993) provides the framework in which the semantics is couched. This is a well motivated approach since versions of channel theory have also been used to provide an analysis of conditional reasoning (Cavedon 1993) (conditional logics themselves having been applied to interpreting inheritance), as well as having been applied directly to the semantics of natural generics (Cavedon and Glasbey 1994) (inheritance reasoning is a model of reasoning with generics).

The components are sites (situations), types, and channels, with relations among the three kinds of objects. $\{T^1, T^2, \ldots\}$ are types; $\{s^1, s^2, \ldots\}$ are sites; $\{c^1, c^2, \ldots\}$ are channels. Let $\{C^1, C^2, \ldots\}$ be types associated with channels. *Signaling relations* pair sites with other sites, relative to a channel ($s^1 \overset{c^1}{\mapsto} s^2$). *Indicating relations* pair types with other types ($T^1 \Rightarrow T^2$). The *of-type* relation pairs sites with types ($s^1 \models T^1$), channels with channel types ($c^1 \models C^1$), and channel types with indicating relations ($C^1 \models (T^1 \Rightarrow T^2)$). It is sometimes easiest to refer to sets of signaling relation 3-tuples in terms of channel types rather than channel tokens ($s^1 \overset{C^1}{\mapsto} s^2$).

Definition 1 *Channel Types. Let $\mathcal{CT}$ be a total function $(c \times C)$ such that $\mathcal{CT}(c') = C'$ iff $c' \models C'$. It is assumed that if $\mathcal{CT}(c') = C'$ and $\mathcal{CT}(c') = C''$ then $C' = C''$.*

I is a function that assigns situation types to nodes in an inheritance network. A model of a default inheritance network is a tuple: $\langle S, T, C, I, \models, \mapsto \rangle$.

Restrictions on the elements of this tuple which obtain models of various inheritance reasoners will be discussed after some essential definitions are spelled out.

The basic idea of this paper is that an inheritance network translates fairly directly into an interpretation involving indicating relations. Each default link in an inheritance network is understood as a constraint in situation theoretic terms. Constraints are interpreted as indicating relations on types of situations and are supported by information channels. Constraints that hold in general may or may not hold in particular. The key to reasoning is composition of the relations named above, but this is only sensible when certain conditions hold with respect to the signaling relations, the actual connections, underneath the indicating relations.

2.1 Tools from Channel Theory

First, it is useful to provide some definitions to constrain models. Some of these come directly from Barwise 1993; others are slightly altered to suit the present purpose or are new. Essentially, these define aspects of unimpeded information flow.

Definition 2 *Given a signaling relation* $s^1 \overset{c}{\mapsto} s^2$, s^1 *is a* signal *for* s^2 *relative to* c *and* s^2 *is a* target *for* s^1 *relative to* c.

A signaling relation is a three-place relation between signal and target tokens and channel tokens. A channel, in this sense corresponds to that of Barwise 1993 and approximately to the sense of *link*[1] used by Barwise and Seligman 1993.

Definition 3 $C \models T^1 \Rightarrow T^2$ *iff for all* $c \models C$, $c \models T^1 \Rightarrow T^2$.

Definition 4 *If* $C \models T^1 \Rightarrow T^2$ *then* T^1 indicates T^2 *relative to* C.

A basic contribution of channel theory is that information channels support constraints. Channels are essentially conduits of information such that if a particular channel is of a certain channel type, then it supports a constraint, and links tokens in such a way that if the signal supports the type in the antecedent of the constraint then the target is classifiable in terms of the consequent of the constraint. Thus, channels license classification of target sites in terms of the consequent types in the constraints supported by the channel type.

Definition 5 *A channel type* $C \models T^1 \Rightarrow T^2$, *is* informative *iff there exists a channel* $c \models C$ *and* s^1 *and* s^2 *such that* $s^1 \overset{c}{\mapsto} s^2$ *with* $s^1 \models T^1$ *and* $s^2 \models T^2$.

[1]This use of *link* is again different from the sense in this paper—that of a link in an inheritance hierarchy. For Barwise and Seligman 1993, *link* denotes certain kinds of connections between classification systems.

Definition 6 *If C is an informative channel type supporting the constraint $T^1 \Rightarrow T^2$, then for all sites s, s' such that $s \overset{c}{\mapsto} s'$, C* classifies its targets *s' as T^2 ($s' \overset{C}{\models} T^2$).*

A part of the modeling power of channel theory comes from the potential difference between the types supported by a token site ($s' \models T^2$) and the types that a token site can be classified in terms of due to the constraints that site is connected to through information channels ($s' \overset{C}{\models} T^2$). In the latter, s' is classified as T^2 by virtue of the constraint/channel type C.

I assume that tokens do not support inconsistent types ($\neg\exists s, s \models T \wedge s \models \neg T$). However, a token can be classified by inconsistent types through different channel types ($s \overset{C}{\models} T \wedge s \overset{C'}{\models} \neg T$). Such a token will be described as having *contested* types. Nothing of substance here hinges on assuming that tokens do not directly support inconsistent types independently of classification via the consequents of different constraints. However, many people seem to feel that inconsistent situations are not possible. For an example of tokens with contested types it is possible to look to sign-based theories of linguistic information (Pollard and Sag 1987,1994) in which *signs* are linguistic types that classify particular utterance situations (see also ideas in situation theoretic grammar (Cooper 1991)). Utterances that are *ungrammatical* but which nonetheless retain semantic interpretability are examples of situations which support inconsistent linguistic types, either directly or via classification supplied by linguistic constraints.[2]

A channel type is sound if all the channels of that type support the classification of targets with types that agree with the types supported by the targets on their own.

Definition 7 *A channel type $C \models T^1 \Rightarrow T^2$, is* sound *iff for all $c \models C$ and for all tokens s^1 and s^2 if $s^1 \overset{c}{\mapsto} s^2$ and $s^1 \models T^1$ then without contest $s^2 \models T^2$.*

Definition 8 *A constraint $T^1 \Rightarrow T^2$ is* sound *iff it is the type of some sound channel, and the constraint is* informative *iff it is the type of some informative channel.*

Seligman and Barwise 1993 characterize channel theory as a theory of information flow. Classification offers the basic unit of information flow: some token is of some type. Channel theory provides a formal means for discussing constraints because it allows description of some situation being of some type carrying information that some other situation is of a particular type by virtue of the connectedness of the two tokens. I distinguish be-

[2] Vogel and Cooper 1995 give some further details about the application of channel theory to feature structures and inconsistent information.

tween sound and informative constraints. Both are types of regularity, but informative constraints are the weaker of the two. A sound constraint is one that classifies a channel whose signal/target pairs in all of the signaling relations it participates in are classified by types consistent with those predicted by the constraint. A regularity, embodied by a constraint, is informative if there is at least one instance of it grounded in a signal/target pair of the predicted types. This condition is stronger still than if the constraint did not exist but two situations (corresponding to the signal/target pair) happened to support the same types (as in a humean correlation).

Defaults are interpreted as informative constraints. Necessarily exceptionless statements (squares have four sides) are interpreted with sound constraints over reflexive channels.[3] Defaults are generally unsound. Recall that a site can be classified by inconsistent types. Such a site can still be part of an informative channel, but not of a sound channel, unless the conflicting types which classify such a site are unrelated to the constraint supported by the channel. Given a particular sentence, *Birds fly* represented with a link *Birds* $\longrightarrow$ *Fliers*, the sentence is interpreted using channels in the following way: there is an indicating relation between situations of a type *Bird* and of the type *flier*; underneath this indicating relation is a signaling relation:

(1)
$$\begin{array}{ccc} Birds & \Longrightarrow & Fliers \\ \amalg & & \amalg \\ s^1 & \overset{c}{\mapsto} & s^2 \end{array}$$

The signaling relation is a connection which is of a channel type that supports the constraint $Birds \Rightarrow Fliers$ and links two situations, one that supports the fact that a bird is present and another that supports the fact that a flier is present. Note that the target site could also support the fact that a flier is not present.

(2)
$$\begin{array}{cccc} Birds & \Longrightarrow & Fliers & \\ \amalg & & C\amalg & \\ s^1 & \overset{c}{\mapsto} & s^2 & \models \neg Fliers \end{array}$$

If I watch a bird eating breadcrumbs in a park, then I am in a situation in which a flier is present; through the information channel, the bird is classifiable as a flier. However, I am also in a situation in which a flier is not present, as the bird is eating breadcrumbs and not flying.[4] The constraint is informative because there are situations which exist in the signaling relation that support the right types, but it is not sound since

[3]A reflexive channel is one that connects a site to itself; these are sometimes also called *redescription* channels, because the site is re-classified through the channel type by the type in the consequent of the corresponding constraint.

[4]It is a flier and not a flier for different reasons, because of different channels.

there are some situations in which birds are present which are not also situations in which fliers are present.

Definition 9 *A full model of a network is informative if all of its constraints are informative, and it is sound if all of its constraints are sound.*

2.2 Defeasibility

Information does not always flow as it seems it should. The idea of channel theory is that channels are still able to model regularity in the world, even when information flow does not correspond to that regularity in particular circumstances. The following definitions are introduced to model the conditions that arise with the interpretation of defaults.

Definition 10 *A channel type* $C \models T^1 \Rightarrow T^2$ *has a* pseudosignal s^1 *iff* $\exists c \models C$ *and* $s^1 \models T^1$ *but there is no* s^2 *such that* $s^1 \overset{c}{\mapsto} s^2$.

Definition 11 *A constraint* $T^1 \Rightarrow T^2$ *has a* pseudosignal s^1 *iff the constraint is the type of a channel c which has* s^1 *as a pseudosignal.*

Given a constraint on situation types, a pseudosignal to that constraint can be understood as an event of the type antecedent to the constraint which is not connected by that constraint to any event which is of the consequent type. Some other event out there may be of the consequent type, but the pseudosignal is not connected to it. For example, Opus the penguin is a pseudosignal for the constraint that *Birds fly*, even if on some particular occasion Opus is flying on a Pan-Am jet. Though the second situation supports the right type, it is not connected via a channel of the type, *Birds fly.* Intuitively, a sound channel has a pseudosignal in any site that is classified by a type that predicts that the site should be a signal, but which is indeed not a signal to any target through the channel. An example of a pseudosignal is given in (3).

(3)
$$\begin{array}{ccc} T^1 & \Longrightarrow & T^2 \\ \Vert & & \\ s^1 & & \end{array}$$

The nature of defaults is that they have exceptions.

Definition 12 *A channel type* $C \models T^1 \Rightarrow T^2$*, has an* exception s^1 *iff there exists* $c \models C$ *and there are tokens* s^1 *and* s^2 *such that* $s^1 \overset{c}{\mapsto} s^2$ *and* $s^1 \models T^1$ *but* $s^2 \not\models T^2$.

An exception is closely related to a pseudosignal in that it is a point of breakdown in classifications. In an exception to a constraint, there is an event that is of the type antecedent to the constraint, and it is connected by a channel that supports the type of the constraint to another event, but that connected event fails to be of the consequent type. For example, Jonathan the seagull that I watched eat breadcrumbs in the park is a signal for the constraint that *Birds fly*, and can be connected via a channel which supports that constraint to a situation in which Jonathan is eating, one which does not support the consequent type of *flight*. Examples of two sorts of exceptions are shown in (4 & 5). In both cases, if the channel type C is otherwise informative then $s^2 \overset{C}{\models} T^2$, even though $s^2 \not\models T^2$. Since any target of a channel is classifiable in terms of the consequent type in the constraint that the channel type supports, it is not necessary to indicate that classification (e.g. $s^2 \overset{C}{\models} T^2$) in diagrams like (4 & 5).

(4)
$$\begin{array}{ccc} T^1 & \Longrightarrow & T^2 \\ \amalg & & \\ s^1 & \overset{c}{\mapsto} & s^2 \end{array}$$

(5)
$$\begin{array}{cccc} T^1 & \Longrightarrow & T^2 & \\ \amalg & & & \\ s^1 & \overset{c}{\mapsto} & s^2 & \models \neg T^2 \end{array}$$

Before proceeding further with the model, it is illuminating to contrast the current approach from that motivated by conditional logics. The main idea behind the conditional logic approach to defaults Delgrande 1987, Delgrande 1988, Delgrande 1990 is to use a modal logic with a conditional operator $\Rightarrow$, where $A \Rightarrow B$ means, "Unless there is an exceptional state of affairs, if A then B." Delgrande develops a first order modal system (which is thus able to represent strict as well as default information using the modal necessity operator for strict relations), and the propositional subset of this logic is equivalent to S4.3. The accessibility relation (E or $\geq$) associated with the logic is reflexive, transitive, and forward connected,[5] where $w_i E w_j$ denotes that w_j is 'at least as unexceptional as' w_i (w_j 'is at least as normal as' w_i). Given a sentence α, $[\![\alpha]\!]^M$ denotes the set of worlds in the model M in which α is true and that set is identified as the proposition given by α. Delgrande provides a *world selection function f* that maps from a particular world and the set of worlds given by a proposition to the set of E-accessible worlds that are at least as unexceptional (most normal). Truth ($\models$) is relative to a world and model.

Definition 13 $f(w, [\![\alpha]\!]^M) =$
$\{w_i | w \geq w_i \;\; \textit{and} \;\; \models^M_{w_i} \alpha, \;\; \textit{and} \; \forall w_j \;\; \textit{such that} \; w_i \geq w_j \;\; \textit{and} \models^M_{w_j} \alpha, w_j \geq w_i\}$

Thus the truth conditions for $A \Rightarrow B$ can be articulated as in Definition 14. This just means that $A \Rightarrow B$ is true at a world w if the worlds at least as normal as w where A is true are a subset of the worlds where B

[5] A relation R is forward connected iff $w_i R w_j$ and $w_i R w_k$ implies either $w_j R w_k$ or $w_k R w_j$.

is true; that is, in the worlds more exceptional than w where A is true, B need not be true.

Definition 14 $\models^M_{w_j} A \Rightarrow B$ *iff* $f(w, [\![A]\!]^M) \subseteq [\![B]\!]^M$.

Definition 15 $\models^M_{w_j} A \supset B$ *iff if* $\models^M_{w_j} A$ *then* $\models^M_{w_j} B$.

Given that the truth conditions for implication are standard (see Definition 15), it would be consistent to have for some constant c, $A(c) \Rightarrow B(c)$ and $A(c) \supset \neg B(c)$, because this just entails that there are no normal worlds where $A(c)$ is true. However, this means that the logic cannot support *modus ponens* for the conditional operator. In the channel theoretic approach, it is not necessary to appeal to a space of possible worlds with a normalcy ordering. Rather, the truth of a default depends just on the existence of an informative constraint holding in the actual world. What the conditional logic captures in the relations between worlds, the current approach captures with a reified notion of constraint. The advantage of this is that the graphic-theoretic nature of inheritance proof theory can be replicated in the semantics, thereby enabling the specification of semantics for a number of inheritance systems.

2.3 Conflicting Defaults

A key feature of path-based default inheritance reasoning is that the meaning of certain links is mitigated by the surrounding network. For instance, exceptions to certain defaults arise from conflicting defaults. The next definitions articulate important relations that can exist between a constraint and other accepted constraints.

Definition 16 *A channel type* $C \models T^i \Rightarrow T^j$, *has a* dual signal s^k *iff* C *is informative and there exists an informative channel type* C' *which supports the type* $T^k \Rightarrow \neg T^j$ *and* $s^k \models T^k$.

Definition 17 *A constraint* $T^i \Rightarrow T^j$ *has a* dual constraint *iff the constraint is the type of a channel* c *which has* s^k *as a dual signal.*

Dual constraints can be understood in classical logic terms as implications which can both be true at the same time only when one of them is vacuously true. For an informative constraint to have a dual is for there to be another informative constraint that carries the opposite information, even though it may be information about different source and target tokens.

A condition stronger than the existence of dual signals is the existence of antisignals. A pair of antisignals is a pair of dual signals that share targets.

Definition 18 *A channel* $c \models T^i \Rightarrow T^j$, *has an* antisignal s^k *iff* s^k *is a dual signal for* C *through* C' *and any target relative to* C *is also a target relative to* C'.

Definition 19 *A constraint $T^i \Rightarrow T^j$ has an antisignal s^k iff the constraint is the type of a channel c which has s^k as an antisignal.*

An antisignal to a constraint is something which conveys contradictory information about target tokens. An antisignal is the source of information which may flow through a completely different sort of channel, information that some token has more than one and conflicting types. An example of a set of classifications that establishes an antisignal is given in (6). Let $c^1, c^2, c^3 \models C$ and $c^4, c^5, c^6 \models C'$; thus, $C \models T^0 \Rightarrow T'$ and $C' \models T^1 \Rightarrow \neg T'$. Both C and C' are informative: C is because $s' \models T'$, and C' is because $s'' \models \neg T'$, so s^4, s^5, s^6 are dual signals to C and s^1, s^2, s^3 are dual signals to C'. Since s', s'' and s''' are targets to both C and C', s^4, s^5, s^6 are antisignals to C and s^1, s^2, s^3 are antisignals to C'. Note that $s', s'', s''' \overset{C}{\models} T'$ and $s', s'', s''' \overset{C'}{\models} \neg T'$.

(6)

T^0	$\Longrightarrow$	T'		T^0	$\Longrightarrow$	T'		T^1	$\Longrightarrow$	$\neg T'$	
$\amalg$		$\amalg$		$\amalg$				$\amalg$		$\amalg$	
s^1	$\overset{c^1}{\mapsto}$	s'		s^3	$\overset{c^3}{\mapsto}$	s'''		s^5	$\overset{c^5}{\mapsto}$	s''	
T^0	$\Longrightarrow$	T'		T^1	$\Longrightarrow$	$\neg T'$		T^1	$\Longrightarrow$	$\neg T'$	
$\amalg$				$\amalg$				$\amalg$			
s^2	$\overset{c^2}{\mapsto}$	s''	$\models \neg T'$	s^4	$\overset{c^4}{\mapsto}$	s'	$\models T'$	s^6	$\overset{c^6}{\mapsto}$	s'''	

Definition 20 *A model is* strict *if it is informative and if none of its channels have antisignals, and it is* defeasible *otherwise.*

2.4 Basic and Composite Channels

Applying channel theory to the semantics of path-based inheritance requires composition of channels to correspond to inheritance paths, or what in a traditional logic corresponds to sanctioning the validity of an inference through a chain of implications. The semantics of default inheritance reasoning, rather than requiring variations of parallel composition, will make crucial use of various serial composition operations. Barwise also provides a definition of serial composition of channels; it is useful to see how this compares with the definition required for the interpretation of inheritance:

Definition 21 *A channel c is the* putative serial composition *of n channels $c^1 \ldots c^n$ $(c^1; \ldots; c^i; \ldots; c^n)$ iff for all sites $s^0, s^n \in S$, $s^0 \overset{c}{\mapsto} s^n$ iff there are $n-1$ intermediate sites (s^i) such that $s^{i-1} \overset{c^i}{\mapsto} s^i$ and $s^i \overset{c^{i+1}}{\mapsto} s^{i+1}$ $(1 \leq i < n)$.*

Barwise actually refers to Definition 21 as serial composition simpliciter (and defines it instead as a binary operator), but given that a general semantics for inheritance requires a number of different forms of composition, it is useful to distinguish putative composites from those actually to be included in a particular model. Definition 21 generalizes to channel types in Definition 22.

Definition 22 *A channel type C is the* putative serial type-composition *of n channel types $C^1 \dots C^n$ ($C^1; \dots; C^n$) iff there exists c' and for all i from 1 to n there exists c^i such that $c^i \models C^i$, where c' is the putative serial composition of $c^1 \dots c^n$.*

Definition 23 *Let $c = (c^1; \dots; c^i; \dots; c^n)$ be the putative serial composition of $n \geq 2$ channels such that for $i = 1$ to n, $c^i \models T^i \Rightarrow T^{i+1}$, then $c \models T^1 \Rightarrow T^n$.*

Given a composite channel $c = (c^1; ...; c^n)$, it is handy to be able to refer to the components of the channel $\{c^1, ..., c^n\}$. In what follows, various versions of serial composition will be more restricted in applicability than putative serial composition and will each have analogous generalizations to composition of channel types.

Consider the structure: $\langle S, T, \mathcal{C}, I, \models, \mapsto \rangle$. Let $\langle S, T', \mathcal{C}', I', \models', \mapsto' \rangle$ be the closure of the former structure under putative serial composition. It is useful to define some simple concepts from these two structures.

Definition 24 *Let the channels in $\mathcal{C}$ be* atomic *or* prime *channels. The composites in $\mathcal{C}' - \mathcal{C}$ are* nonatomic.

Definition 25 *Let c be a channel; PRIMES(c) denotes the set of atomic channels that comprise c.*

Definition 26 *Let c and c' be channels; c and c' are* site-equivalent *iff $\forall s^1, s^2$, $s^1 \overset{c}{\mapsto} s^2$ if and only if $s^1 \overset{c'}{\mapsto} s^2$.*

Definition 27 *An atomic channel $c \models T \Rightarrow T'$ is* basic *iff for any c' which is a nonatomic site-equivalent to c, c' and c support mutually inconsistent constraints.*

Consider (7) with $c^1, c^2 \models C$; $c^3, c^4 \models C'$; $c^5, c^6 \models C''$.

(7)

$\begin{array}{cccc} T' & \Longrightarrow & \neg T'' & \\ \amalg & & & \\ s' & \overset{c^3}{\mapsto} & s^2 & \models T'' \end{array}$	$\begin{array}{ccc} T & \Longrightarrow & T' \\ \amalg & & \amalg \\ s^1 & \overset{c^1}{\mapsto} & s' \end{array}$	$\begin{array}{ccc} T & \Longrightarrow & T'' \\ \amalg & & \amalg \\ s^7 & \overset{c^7}{\mapsto} & s^8 \end{array}$
$\begin{array}{cccc} T & \Longrightarrow & T' & \\ \amalg & & & \\ s^3 & \overset{c^2}{\mapsto} & s'' & \models \neg T' \end{array}$	$\begin{array}{ccc} T' & \Longrightarrow & \neg T'' \\ \amalg & & \amalg \\ s'' & \overset{c^4}{\mapsto} & s^4 \end{array}$	$\begin{array}{ccc} T & \Longrightarrow & \neg T'' \\ \amalg & & \\ s^3 & \overset{c^5}{\mapsto} & s^4 \end{array}$
	$\begin{array}{ccc} T & \Longrightarrow & \neg T'' \\ \amalg & & \amalg \\ s^5 & \overset{c^6}{\mapsto} & s^6 \end{array}$	

Let $c' = (c^1; c^3)$, $c'' = (c^2; c^4)$; then, $c', c'' \models T \Rightarrow \neg T''$ (the same as c^5 and c^6). Also, $C^{\triangle} = (C; C')$. The atomic channels are $\{c^1, c^2, c^3, c^4, c^5, c^6\}$, and c', c'' are nonatomic. $PRIMES(c') = \{c^1, c^3\}$; $PRIMES(c'') = \{c^2, c^4\}$. The nonatomic channel c'' is site-equivalent to the atomic channel c^5 but not to c^6. Therefore, c^6 is basic but c^5 is not (and $c^1 \ldots c^4$ are all basic). Note that C'' and $C^{\triangle}$ both have a dual signal in s^7 through C'''.

Also consider a notion that can be defined in channel theoretic terms as *originality.* A channel is original if it is basic, or if it is not basic it is original if it supports the flow of information that is not supported by its constituent channels.

Definition 28 *An atomic channel* $c \models T \Rightarrow T'$ *is* original *iff*

1. *c is basic, or*
2. *c has a nonatomic equivalent decomposable into* $\{c^1 \ldots c^n\}$ *such that for some* $i, 1 \leq i \leq n$, $c^i \models C$ *and* C *has a dual signal.*

In (7), although c^5 is not basic it has a nonatomic equivalent c'' which has a dual signal through c^7 in C'''. Thus, c^5 is original. The channels $c^1 \ldots c^4, c^6$ are all original because they are basic. The channel c^7 is basic as well.

Definition 29 *An atomic channel type* C *is* original *if all* c *such that* $c \models C$ *are original.*

Definition 30 *A channel* $c \models T \Rightarrow T'$ *is* effectively original *iff*

1. *c is original, or*
2. *c is decomposable into atomic channels* $c^1 \ldots c^n$ *such that each* c^i *is original.*

Definition 31 *A channel type* C *is* effectively original *iff each* c *such that* $c \models C$ *is effectively original.*

Effective originality extends the notion of original contributions of information from primes to putative composites. In the definition of composition offered below, composition with effectively unoriginal channels is disallowed. This is sensible, because, by definition, the same information is already conveyed elsewhere if a putative composite is effectively unoriginal.

Now it is possible define the directness of a channel (denoted, $|c|$) as the number of original channels in its putative composition; if c is a basic channel then $|c| = 1$. This admits the possibility of organizing channels into a specificity hierarchy. The one most useful for interpreting path permission in H90[6] follow:

[6] This is the reasoner presented by Horty et al. 1990. Vogel et al. 1993 present the definition of the system using the declarative specification assumed here. Vogel 1995 parameterizes this definition to realize other reasoners from the literature that are in the same family.

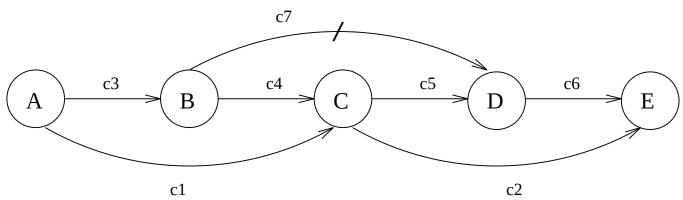

Figure 1 A Network with 7 Atomic Links, 5 Basic Links, 6 Original Links.

Definition 32 *Given two channels c^1 and c^2, each of the type $T^1 \Rightarrow T^2$ or $T^1 \Rightarrow \neg T^2$, c^1 is at least as direct as c^2 ($c^1 \preceq c^2$) iff $|c^1| \leq |c^2|$.*

Thus, one channel is more direct than another if it connects the same points and is composed from fewer original channels.

It is easiest to understand the definitions of basic channels and of original channels given above, in terms of the network topologies that they interpret. Consider the graph from Figure 1. The network in this figure is labeled with channel types that interpret those links. The links *ac* and *ce* are interpreted by nonbasic channel types, and the rest are basic. The link $b \not\longrightarrow d$ is interpreted as original since the node *c* offers a dual signal to c^2, and the link *ce* is original since *b* yields a dual signal for c^5, but the link *ac* is not original.

3 A Model of Inheritance Reasoning

Let Γ be an inheritance network. Let M be a tuple: $\langle S, T, \mathcal{C}, I, \models, \mapsto \rangle$; S is a nonempty set of sites and T is a set of types. I is an interpretation function which assigns a unique type to each node in the inheritance network. For each link $n_1 \longrightarrow n_2$ in Γ there is an informative channel type $C \in \mathcal{C}$ of the type $I(n_1) \Rightarrow I(n_2)$, and for each link $n_1 \not\longrightarrow n_2$ there is exactly one informative channel type C that supports the constraint type $I(n_1) \Rightarrow \neg I(n_2)$. Note that there are no constraints of the form: $\neg\ I(n) \Rightarrow \tau$ for any node n or any type τ. For each channel type $C, C' \in \mathcal{C}$, if C supports a constraint whose consequent type is identical to the antecedent of a constraint supported by C', then for some $c \models C$ and $c' \models C'$ there exists a site s which is a target to C and a source to C'. There are no other channels in M. A model of inheritance reasoning over Γ is given by M', a tuple $\langle S', T', \mathcal{C}', I', \models', \mapsto' \rangle$ derived from M via the closure under the appropriate version of serial composition of channels (to be defined in the course of this section), with the assumption that no target to a channel is also a target to a channel composed from it. A path π through Γ is supported by a particular reasoner if and only if $I(\mathit{firstnode}(\pi)) \Rightarrow I(\mathit{lastnode}(\pi))$ is a constraint supported by an informative channel type $C \in \mathcal{C}$ in the corresponding closure. Given an inheritance network Γ, π is a positive path in Γ iff there exists an informative channel C such that $C \models$ *I*(*firstnode*(π)) $\Rightarrow$ *I*(*lastnode*(π)) in M', and π is a negative path if and only if there exists an informative channel C such that $C \models$ *I*(*firstnode*(π)) $\Rightarrow \neg$ *I*(*lastnode*(π)) in M'.

3.1 Interpreting Permitted Paths: Restricted Skepticism, Off-Path Preemption

Just as there is assumed be a channel theoretic constraint underlying each link in an inheritance network, there is also assumed to be a composite constraint underlying each of the permitted paths in a network. Since permitted paths are defined by the exact specifications of the reasoning algorithm, this means that the channel theoretic interpretation of an inheritance network also provides an interpretation to the exact reasoning mechanism applied to that network. For instance, the definition of channel composition that is given in Definition 33 guarantees that composite channels exist for only those paths that are not preempted (by the H90 reasoner).

Finally, a definition of serial composition takes these notions into account. The serial composition of two channels can be formed as a channel if an intermediate site exists that supports the right types and if no more direct channel spans the two composed channels forming an antisignal to their composition; since this definition is particular to a model of H90 it is called *H90 serial composition.*

Definition 33 *A channel* $c \models T \Rightarrow T'$ *is the* H90 serial composition *of* $c^1 \models T \Rightarrow T''$ *and* $c^2 \models T'' \Rightarrow T'$ *(*$c^1;c^2$*) iff*

1. *c is the putative composition of* c^1 *and* c^2 *where* c^1 *is an original channel or an effectively original H90 serial composite and* c^2 *is original, and*
2. *for any antisignal s through a putative channel* c^i *there exists a channel* c^j *through which s is a antisignal and a channel* $c^\dagger$ *supporting the same constraint as c, such that the following conditions are satisfied:*
   a. *PRIMES(*c^i*)* $\cap$ *PRIMES(*c^j*)* $\neq \emptyset$
   b. *PRIMES(*c^j*)* $\cap$ *PRIMES(*$c^\dagger$*)* $\neq \emptyset$
   c. $c^\dagger \prec c^j$

The idea behind this restricted serial composition is that it is the complete interpretation of off-path preemption in H90. Essentially, there is no composite channel underneath paths composed from redundant links. Redundant links are unoriginal channels. Note that the $c\dagger$ mentioned in this definition could in fact be the putative composite c itself. If this occurs in the model of an inheritance network, then it is an instance of on-path preemption. Allowing $c\dagger$ to be distinct from c in its composition provides an interpretation to off-path reasoning.

For a given network Γ, assume the existence of M and M' as constructed before. Let M^{H90} be the closure of M under H90 serial composition; M^{H90} constitutes a model of reasoning in H90 over Γ. That is, if π is a positive path in Γ permitted by H90, then $c : I(\mathit{firstnode}(\pi)) \Rightarrow I(\mathit{lastnode}(\pi))$ is an informative channel in M^{H90}, and if π is a negative path in Γ permitted by H90, then $c : I(\mathit{firstnode}(\pi)) \Rightarrow \neg I(\mathit{lastnode}(\pi))$ is an informative channel in M^{H90}. Moreover, if $c : \phi \Rightarrow \psi$ is an informative channel in M^{H90} then a path π in Γ supports that conclusion under the reasoning definitions of H90.

H90 serial composition is a binary operator (just as is path construction in H90), which assumes that the first channel type is either original, or an H90 composite (thus, a channel type from M^{H90} and not just M'), and that the second channel type is original. This entails that if a channel type C is in M^{H90} (and not in M) it corresponds to a path in Γ composed of a permitted path extended by a single link where the resulting path is neither preempted nor conflicts with any path. Proposition 1 holds.

Proposition 1 *Given an inheritance network* Γ, π *is a positive path in* Γ *if and only if there exists an informative channel* C *such that*
$C \models \mathrm{I}(\mathrm{firstnode}(\pi)) \Rightarrow \mathrm{I}(\mathrm{lastnode}(\pi))$ *in* M^{H90}, *and* π *is a negative path if and only if there exists an informative channel* C *such that*
$C \models \mathrm{I}(\mathrm{firstnode}(\pi)) \Rightarrow \neg\, \mathrm{I}(\mathrm{lastnode}(\pi))$ *in* M^{H90}.

Modifications to the basic definition which provides an interpretation to the H90 system can be provided to create interpretations for the other possible modes of reasoning. The H90 system is restrictedly skeptical (hence skeptical and not credulous), utilizing off-path preemption, and exhibiting stability when faced with topologically redundant links.

3.2 Inferring Properties of Individuals

A natural channel theoretic interpretation of inference in classical non-modal propositional logic requires only a single token. Types correspond to sentences in the logic. Conclusions from inferences can be looked up in the semantics in the types assigned to a token. The sentence $A \supset B$ is true iff there is a constraint $A \Rightarrow B$ that is the type of an identity channel on the single token. If the channel is sound, then if the token is of type A then the token will also be classified as B. This is essentially the idea of a *redescription* channel.

As has been said, inheritance links are interpreted as informative constraints, and the existence of a channel type that supports a constraint is sufficient to classify target tokens relative to the channel type with the type that is consequent to the constraint. The channels that underlie the constraints generally link distinct tokens. Information that some token is one way carries information about another token being some other way.

Of course, token identity is possible, but the framework admits much more into its descriptive auspices. On the other hand, inference is performed using inheritance as an efficient way of calculating properties of some object at hand[7], and other objects whose properties are conveyed by constraints are not entirely relevant to an inference, although the properties possessed by those objects are relevant. Inference in default inheritance networks is aptly interpreted by the projection of types related to tokens via constraints (the types of targets) upon signaling tokens, analogously to the way an informative channel licenses a particular classification of its target tokens. Thus, inference using the constraint *Birds fly* when applied to the token *Opus*, which is a bird (and in the absence of other constraints), leads to the classification of *Opus* as a *Flier*. This is distinct from a redescription channel in which signal and target are actually the same token. The room for error in projecting classifications accounts for the classical invalidity of default inheritance, but its reasonableness, apart from compact representation is also arguable from the similarity to the model of inference in the classical case. The purpose of inference using defaults is to avoid having to look up information in a deductive table. In this case it would mean looking up the *Opus* token in that related classification to see if *Opus* is flying in there, and similarly for any other object that one could try to reason about.

Definition 34 *Let* $\chi = \langle S, T, C, \rightarrow, \Rightarrow, \models \rangle$ *be a system of constraints on a classification domain. The* leftward projection *of an informative constraint* $T^1 \Rightarrow T^2$ *down a channel* c *of type* C *such that* $C \models T^1 \Rightarrow T^2$ *upon a specific token* s^o *that is a signal for the constraint yields a new classification and system of constraints,* χ' *such that:*

1. $S' = S, T' = T, C' = C, \Rightarrow' = \Rightarrow$
2. $\rightarrow' = \rightarrow$ *except for those elements of the signaling relation that use the channel* c*. Replace each triple* $\langle s^o, c, s^i \rangle$ *in* $\rightarrow$ *with* $\langle s^o, c, s^o \rangle$ *in* $\rightarrow'$*.*
3. $\models' = \models$ *except for those elements of the of-type relation that use sites and types salient to the constraint projected upon the source. Specifically, replace each* $\langle s^i, T^j \rangle$ *such that* $\langle c, T^h \Rightarrow T^j \rangle$ *is in* $\models$ *and* $\langle s^o, c, s^i \rangle$ *is in* $\rightarrow$*, with* $\langle s^o, T^j \rangle$*.*

The leftward projection given in Definition 34 formalizes the notion described above (see Seligman and Barwise 1993 for alternative forms of projections). It is focused upon a specific token s^o about which reasoning is performed. All of the types, tokens, channels, and constraints from the original classification remain in the projection, but the signaling and of-type relations change. The change in the of-type relation is the essence

[7]Assuredly, there are plenty of inefficient inheritance reasoners available, but H90 is polynomial.

of the projection: types supported by a target relative to the channel the constraint is projected along are attributed by the source instead. Since the constraint and channel remain, the leftward projection also redirects the signaling relation so that the former target is no longer signaled, with the source signaled instead.

4 Conclusions

Channel theory provides a formal framework which admits distinctions of fine granularity attuned to the needs of providing semantics to the family of inheritance reasoners under consideration. The sanctioning of a path in a reasoner is interpreted by the existence of an informative constraint supported by either an atomic or composite channel in the corresponding closure of the atomic model. The restrictions on path permission imposed by particular inheritance proof theories are modeled by restrictions on the composition of channels. In using inheritance reasoning to ask if something represented by some node has a property represented by some other node the question is posed in terms of the existence of a permitted path between the two nodes. If there is a path and it is positive, the object represented by the first node is assumed to have the property represented by the second node, and if the path is negative, then the object is considered to have the antiproperty. In the interpretation provided here, the existence of a path between the two nodes is interpreted as by the existence of a natural regularity of the denoted kind. That is distinct from the question of whether the object being considered does indeed behave in accordance with that natural regularity. It could still be an exceptional object of some sort. To answer questions about objects in the semantics, it is necessary to simply look at the types that classify it. In the case that certain types don't classify it explicitly then it is possible to perform a leftwards projection of the types that the object participates with in constraints (see §3.2). If leftwards projection is used in conjunction with H90 serial composition, the conclusions of skeptical inference about individuals are available.

Conditional logics have been used in the past to give semantics to individual inheritance reasoning systems (Delgrande 1988, Boutilier 1989) (and problems in conditional reasoning has also been analyzed via channel theory (Cavedon 1995)). An inheritance network is translated into the logic, mapping default links to corresponding sentences that use a conditional operator. The possible worlds semantics underlying the logic thus serves as the semantics for reasoning over the network. The essence of the semantics is the normalcy relations among possible worlds which supply the truth conditions for the conditional operator. The result is that a given logic can provide a semantics to only one single form of inheritance reasoning; it isn't possible in this way to capture the semantics of a family of systems.

The channel theoretic approach given here contrasts from the conditional logic approach in finding a direct interpretation from links in a network to objects in the world. A default is true if there is a corresponding informative constraint in the world, and these are modeled directly as objects in the world as described above. Reasoning over the networks corresponds to composing channels. It is the fact that various kinds of restrictions can be placed directly on the allowable compositions that allows this approach to easily offer semantics to a range of inheritance reasoners, and not just the one that agrees with the stated semantics. The graph theoretic specification of path preference in an inheritance reasoner is interpreted directly by the constraints on channel composition.

The channel theoretic framework provides a general semantics for path-based inheritance reasoning, because the different conditions on channel composition interpret the path permission relations of various members in a large family of path-based inheritance reasoners (Vogel 1995). While this work focuses on the semantics of defeasible inheritance reasoning, the approach is general enough to provide semantics to inheritance with strict links as well (see Horty and Thomason 1988); in that case, strict links are modeled by redescription channels rather than channels linking distinct tokens. Past work has tended to either provide a semantics to a particular system (Touretzky 1986), or to work from a semantic framework to a particular inheritance logic whose proof theory approximates reasoning in the desired models (Delgrande 1988, Boutilier 1989, Veltman 1994). Other work in giving a semantics to inheritance reasoning explains proof theoretic differences in terms of different "process models" without actually stating what the different process models are (cf. Dimopoulos 1992). By stipulating the actual conditions under which channels may compose, this work can be seen as giving a "process model". However, the basic principle is that channels are objects that offer exactly the right level of description for giving a general semantics to path-based inheritance reasoning.

Acknowledgements

Thanks to Robin Cooper, Lawrence Cavedon and Jon Barwise for detailed comments, suggestions and discussion about this material. The suggestions of an anonymous reviewer have improved this paper. Remaining faults are my own. I am grateful to the Marshall Aid Commemoration Commission for administering the Marshall Scholarship which made it possible for me to do a Ph.D. in Edinburgh, and to Hans Kamp for an amazing post doc in Stuttgart.

References

Barwise, Jon. 1993. Constraints, Channels, and the Flow of Information. In *Situation Theory and its Applications, Volume 3*, ed. Peter Aczel, David Israel, Yasuhiro Katagiri, and Stanley Peters. 3–27. CSLI Lecture Notes no. 37.

Barwise, Jon, and Jerry Seligman. 1993. Imperfect Information Flow. In *Proceedings of the Eighth IEEE Symposium on Logic in Computer Science*. Washington, D.C., IEEE Computer Society.

Barwise, Jon, and Jerry Seligman. 1994. The Rights and Wrongs of Natural Regularity. In *Philosophical Perspectives, Vol 8*, ed. J. Tomberlin. 331–65. California: Ridgeview.

Boutilier, Craig. 1989. A Semantical Approach to Stable Inheritance Reasoning. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1134–39. Detroit, Michigan.

Brachman, Ronald J. 1983. What ISA Is and Isn't: An Analysis of Taxonomic Links in Semantic Networks. *IEEE Computer* 16(10):30–6.

Cavedon, Lawrence. 1993. Reasoning about Events: The Qualification Problem. Unpublished manuscript. Centre for Cognitive Science, University of Edinburgh.

Cavedon, Lawrence. 1995. *A Channel Theoretic Approach to Conditional Reasoning*. Doctoral dissertation, Centre for Cognitive Science, University of Edinburgh.

Cavedon, Lawrence, and Sheila Glasbey. 1994. An Information-Flow Model of Generics. Presented at the Fifth Symposium on Logic and Language, Nosvaj, Hungary, September 2–6 1994. Submitted to *Acta Linguistica*.

Cooper, Robin. 1991. Three Lectures on Situation Theoretic Grammar. In *Natural Language Processing, EAIA '90, 2nd Advanced School in Artificial Intelligence.*, ed. Miguel Filgueiras, Luis Damas, Nelma Moreira, and Ana Paula Tomás. 102–140. Lecture Notes in Artificial Intelligence, 476. Berlin: Springer Verlag.

Delgrande, Jim. 1987. A First-Order Conditional Logic for Prototypical Properties. *Artificial Intelligence* 33(1):105–30.

Delgrande, Jim. 1988. An Approach to Default Reasoning Based on a First-Order Conditional Logic: Revised Report. *Artificial Intelligence* 36(1):63–90.

Delgrande, Jim. 1990. A Semantics for a Class of Inheritance Networks. In *Proceedings of the 8th Biennial Conference of the Canadian Society for Computational Studies of Intelligence*. Ottawa, Ontario.

Dimopoulos, Yannis. 1992. On the Semantics of Inheritance Networks. *Journal of Experimental and Theoretical Artificial Intelligence* 4(1):19–28.

Horty, John, and Richmond Thomason. 1988. Mixing Strict and Defeasible Inheritance. In *Seventh National Conference on AI*, 427–32. AAAI-88, St Paul, Minnesota.

Horty, John, Richmond Thomason, and David Touretzky. 1990. A Skeptical Theory of Inheritance in Nonmonotonic Semantic Networks. *Artificial Intelligence* 42(2-3):311–48.

Pollard, Carl, and Ivan Sag. 1987. *Information-Based Syntax and Semantics, Volume 1: Fundamentals*. Centre for the Study of Language and Information, Stanford University, CA.

Pollard, Carl, and Ivan Sag. 1994. *Head-Driven Phrase Structure Grammar*. University of Chicago Press and CSLI Publications.

Seligman, Jerry. 1990. Perspectives in Situation Theory. In *Situation Theory and its Applications, Volume I*, ed. Robin Cooper, Kuniaki Mukai, and John Perry. 147–191. Centre for the Study of Language and Information, Stanford University, CA.

Seligman, Jerry, and Jon Barwise. 1993. Channel Theory: Toward a Mathematics of Imperfect Information Flow. Unpublished manuscript.

Touretzky, David. 1986. *The Mathematics of Inheritance Systems*. Los Altos, CA: Morgan Kaufman.

Touretzky, David, John Horty, and Richmond Thomason. 1987. A Clash of Intuitions: The Current State of Non-Monotonic Inheritance Systems. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, 476–82. Milan, Italy.

Veltman, Frank. 1994. Defaults in Update Semantics. Unpublished ms.; Institute for Language, Logic and Information, University of Amsterdam.

Vogel, Carl, and Robin Cooper. 1995. Robust Chart Parsing with Mildly Inconsistent Feature Structures. In *Nonclassical Feature Systems*. Edinburgh Working Papers in Cognitive Science, Andreas Schöter and Carl Vogel (eds.), Volume 10. EUCCS-WP10. pp. 197-216.

Vogel, Carl, Fred Popowich, and Nick Cercone. 1993. Logic-Based Inheritance Reasoning. In *Proceedings of the 9th Bi-Annual Conference of the Society for the Study of Artificial Intelligence and the Simulation of Behavior*, 179–88. AISB'93, Birmingham, England. March 29–April 2, 1993. Appears in, A. Sloman, et al. (eds.), *Prospects for Artificial Intelligence*. Amsterdam: IOS Press. 1993.

Vogel, Carl M. 1995. *Inheritance Reasoning: Psychological Plausibility, Proof Theory and Semantics*. Doctoral dissertation, Centre for Cognitive Science, University of Edinburgh.