Multiple Contradictory Causal Fields

Jeremy J. Carroll, Maria Chiara Menchini*
Telecom Management Department
HP Laboratories Bristol
HPL-97-129
October, 1997

E-mail: jjc@hpl.hp.com

causality,
causation,
paraconsistency,
logic, inconsistent,
networking,
ISO/OSI 7 layer
model

This paper notes that Mackie's concept of causal field can be extended to indicate the domain of acceptable explanations for an event. When we have multiple possible causal fields then it is sometimes desirable to reason using more than one causal field, and there is no *a priori* reason that they should be consistent. We explore this idea by means of two examples; one being a detailed account of the interactions between the layers of the OSI/ISO stack used in computer networking.

*Department of Computer Science, Manchester University, Manchester, UK.

1. Introduction

Recent accounts of causation usually embed causal reasoning within non-standard logics. For example: Lewis [5] uses counterfactual implication with a possible worlds semantics; Shoham [10] uses a modal logic, Reiter [8] uses default reasoning; Menchini [7] uses abductive logics.

This paper looks at one aspect of causation, causal fields, introduced by Mackie [6], and proposes that a paraconsistent[1] treatment may be appropriate. We explore this proposal by examples from medical ethics and computer networking. Towards the end of the paper we examine issues of identity that arise with each of our examples.

2. Causal fields

In causal reasoning Mackie [6] argues that the causal field is a key concept. For Mackie the causal field is a refinement of the question: 'What caused this?' to 'What caused this in these cases but not these others.' For example, if we ask 'What caused this man's skin cancer?' 'exposure to radiation' may be a permissible answer, but not if the causal field is the class of men exposed to radiation. Menchini [7] modifies the concept of causal field to be the set of permissible explanations; 'which of these *a priori* plausible or interesting explanations caused this?' She argues: "explanations provided for a child and for a specialist must satisfy different requirements." She continues by viewing explanations as hypotheses that are abduced as required. In this paper, we follow Menchini's argument one step further and view the causal field not only as the permissible explanations, but also as the set of rules (causal rules, implications, background knowledge, etc.) that is to be used in making an explanation. For example the explanation of sickness in a patient will depend on whether we seek a legal cause (e.g. negligence, which according to a legal judgement is culpable), a medical cause (e.g. a poisoning, which can be shown using medical techniques to have been effective), or a religious explanation (e.g. a punishment for sin, shown using a religious argument to give rise to the sickness).

3. Multiple inconsistent causal fields

We believe that in many real-world reasoning problems it is necessary not only to consider these multiple (potentially inconsistent) causal fields independently, but also together as a single paraconsistent system. We hence propose causal reasoning as a domain of application of paraconsistency.

[1] We use the terms paraconsistent and paracomplete as in the World Congress on Paraconsistency. A paraconsistent logic is a non-trivial logic in which some proposition is true, and its negation is true. A paracomplete logic is one in which some proposition is false and its negation is false.

A doctor in a litigious and religious society must consider the medical, legal, and spiritual dimensions of a proposed treatment. A patient's own sense of their interest may be better served by a doctor working with causal systems including rules:

not Blood Transfusion $\Rightarrow$ Death (Medical)

Blood Transfusion $\Rightarrow$ not Death (Medical)

Death $\Rightarrow$ Liability (Legal)

Blood Transfusion $\Rightarrow$ Eternal damnation (Patient's religious belief) [2]

Eternal damnation $\Rightarrow$ Death (Patient's religious belief)

The different causal fields are at odds with one another. All these rules together would lead to the patient's certain death, and the doctor's liability (but damnation is avoidable). The reality of the doctor's decision is that it depends on all three causal fields together, but not in an undifferentiated way. One can argue that the apparent contradiction between a blood transfusion medically saving the patient and religiously being the death of them as being a simple muddle of terms. We may try to clarify the arguments by distinguishing 'the death of the body' from the 'death of the soul'. Such a move, will, however, be at the expense of religious seriousness, and will misrepresent the deliberate force with which religious beliefs are held.

4. The ISO/OSI networking model

We take further examples of causal fields from the ISO seven layer Open Systems Interconnection networking model, see Tanenbaum [11]:

Layer 7: Application
Layer 6: Presentation
Layer 5: Session
Layer 4: Transport
Layer 3: Network
Layer 2: Data-link
Layer 1: Physical

This is used by system designers, software and hardware engineers, for building computer networks (such as the internet). The underlying idea is that each module of the networking software implements some level or other, and it only concerns itself with that level, and its interactions with the immediately superior and immediately inferior levels. Each software module assumes that the layer below it functions in the way in which it is specified to function. An engineer, building a component of a networking

[2] This rule is not intended to be an accurate reflection of any group's religious beliefs.

system, concentrates only on their layer, and does not have to think in the terms and models of the layers above or below.

So the individual layers provide a simple frame of reference when that is all that is needed, but, a more thorough understanding is reached by reasoning in multiple layers. The ability to do so may be the distinguishing mark of a 'guru' from a mere 'expert'.

Moreover, it is at least conceivable that networking gurus reason in multiple layers using paraconsistent thinking rather than traditional logic.

We will explore this issue by looking at e-mail.

4.1 Sending an e-mail

When Alice sends an e-mail to Bob over the internet her e-mail application (layer 7) may guarantee to either deliver the e-mail, or tell Alice about a problem. For the sake of simplicity we will omit layers 2 to 6, and jump straight to the physical layer. The e-mail will be transmitted by a number of different physical devices, e.g. a 10 baseT Local Area Network, a copper telephone cable, and a transatlantic fibre-optic cable. Each physical link will be responsible for one hop in the journey of the e-mail.

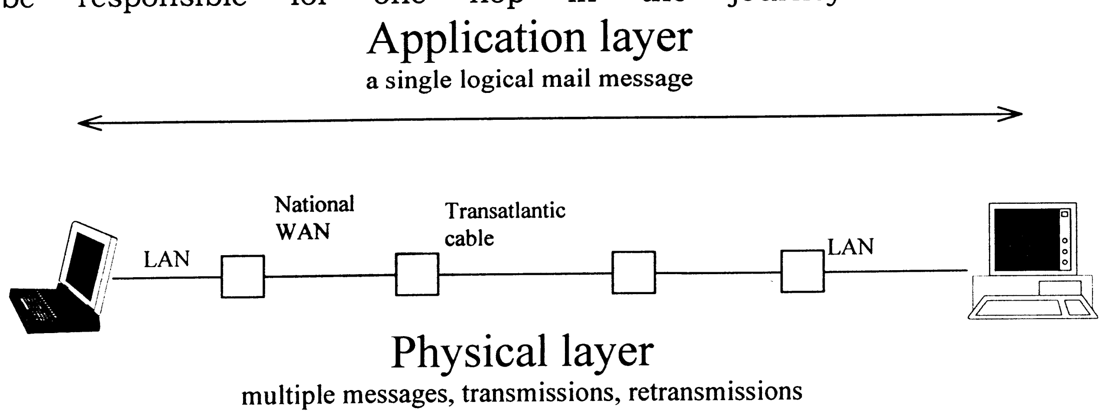

A physical link might break, for example, a ship may snag an underwater cable. If this happens while the e-mail is in transit, then from the perspective of the physical layer, the e-mail is not received. However, at layer 7, the e-mail application will be notified of this failure, and will, at a later time, attempt to retransmit the e-mail. Normally, such retransmission is eventually successful. Thus from the perspective of layer 1 the (original) message is not delivered, while at layer 7 the message is received.

A conventional understanding of this situation will distinguish the message and its delivery in layer 1 from the message and its delivery in layer 7. These are seen as distinct events, (they do occur at different times and in different places).

An alternative understanding would prioritise the unicity of the message; (the same information is indeed being passed at layer 1 and layer 7). Thus, the logical message being passed at the application layer would be identified with the original message(s) at the physical layer (which do not arrive), and would also be identified with the retransmitted message(s), which do arrive. In this identification we would want to conclude that the logical message arrives. We also need to avoid identifying the original transmission at the physical layer, with the retransmission, despite the identification of the logical message with both.

This alternative understanding involves contradiction: the message does not arrive, the message arrives. The physical transmission is the logical message, the physical retransmission is the logical message but the transmission is not the retransmission. Rather than viewing these contradictions as a critique of this understanding, we can take it as a critique of the choice of logic. This motivates moving to a paraconsistent framework that will accept the contradiction without collapsing into triviality.

Such an understanding may be closer to the naive user's common-sense understanding of internetworking; and may also be closer to the guru's mastery of multiple layers. It is almost certainly anathema to the expert's thorough understanding of part of the picture.

The paraconsistent view has the additional advantage of leaving room for an emotional response to the effectiveness of the internet in providing reliable communication over unreliable links, and, with systems such as pretty good privacy [12], in providing private communication over public networks.

4.2 Towards a paraconsistent formulation of networking

We will sketch a formalisation of a paraconsistent reading of the ISO/OSI networking model. This sketch is in terms of Rescher and Brandom's Logic of Inconsistency [9]. They introduce the operators $\dot{\cap}$ and $\dot{\cup}$ over worlds in which a classical logic is interpreted. A formula P is true in the paraconsistent[3] world $w_1 \dot{\cup} w_2$ iff P is true in w_1 or in w_2. Similarly P is true in the paracomplete[4] world $w_1 \dot{\cap} w_2$ iff P is true in w_1 and in w_2.

Given a formula that refers to a message m, this formula can be evaluated in a classical way with m being instantiated as one of the messages in the conventional understanding of the network stack. For example, we could instantiate m as the logical message at the application layer, or as the physical message in the original transmission, or as the physical message in the retransmission (if any). We can then evaluate the formula classically. The formula is paraconsistently true if it is true for any instantiation of the message. The formula is paracompletely true if it is true for all instantiations of the message.

---

[3] They call it an 'inconsistent' world, we use 'paraconsistent' to be in accord with the terminology used at the World Congress on Paraconsistency.

[4] They call it 'schematic, we use 'paracomplete' as in the WCP.

Thus, in the previous example of the e-mail, paraconsistently "the message is received", and "the message is not received". However, it is false that "the message is and is not received", since this last sentence is not true about any message.

A more thorough exposition of this formulation would need to ensure that items other than the message are instantiated appropriately for each instantiation of the message. Using such an instantiation we can evaluate all terms within a formula in a collective way (to use Rescher and Brandom's terminology).

4.3 Causality in the networking stack

In the domain of distributed computing, Birman & Joseph [1] and Carroll & Borshchev [2] propose 'causally' ordered message passing (they differ as to whether a sufficient or necessary condition for causation is more important). They all agree that the causation that should be considered is at the application layer. The problem they address, is that of the non-deterministic order of arrival of messages at a computer in a distributed system. The order of message arrival depends largely on factors in the lower layers of the stack. For the application programmer it may be more convenient if the order in which messages are delivered to the application layer at the computer depends on application layer semantics. Hence, the event of a message arriving (i.e. being passed from layer 5 to layer 6) is separated from the event of the message being delivered (i.e. being passed from layer 6 to layer 7). In particular, the presentation layer may delay delivery of a message.[5] A message is delayed whenever there is still the possibility of other messages arriving that may have to be delivered earlier. (Given any two messages the presentation layer can tell which one should be delivered first.)

Let us consider three computers *A*, *B* and *C*. *B* sends message *b* to *A* and *C* sends message *c* to *A*. If message *b* arrives first, then it is delayed (not delivered) by the presentation layer in *A*. Message *c* arrives; the presentation layer compares *b* and *c*; it decides that *b* should be delivered first; *b* is then delivered.

For the application layer, it is the sending of *b* that causes the delivery of *b*. For the presentation layer, it is the arrival of *c* that causes the delivery of *b*. In the terms of Menchini [7], Birman & Joseph and Carroll & Borshchev take the causal field to be the set of application layer events.

The layers provide a frame of reference for conflicting yet valid causal reasoning, and that when designing, testing and debugging the distributed system we may need to think using multiple causal fields simultaneously.

---

[5] A message is delayed by storing the message in the presentation layer, without immediately passing it to the application layer. Later, it can be delivered from the memory in which it is stored to the application layer.

As before, a conventional approach will draw distinctions between: the message being sent by the application layer of *B*; the message arriving in the presentation layer of *A* and the message being delivered from the presentation layer of *A* to the application layer. In this way, consistency is salvaged at the cost of a loss of message identity.

In contrast, a paraconsistent approach would combine all the layers of the network stack (i.e. all the casual fields). In this way we would see both that the application layer sending of the message *b* is the (unique) cause of the delivery of *b* to *A*, and the presentation layer arrival of the message *c* is the (unique) cause of the delivery of *b* to *A* (with a distributive 'and' rather than a collective 'and'). This view is in contrast to a traditional causal analysis that forces us to prioritise one or other of the two causes and regard the other as merely a background condition.

5. Identity, causation and paraconsistency

We have seen the topic of identity in each of our examples:

- Should we identify the person with the body and the soul, and hence the death of the body with the death of the soul?
- Should we identify an e-mail at the application layer with its counterpart in the physical layer?
- Should we identify a message being sent by computer *B* with the message arriving at computer *A*?

This topic arises naturally in causation. If, like Hume [4], we see causation as repeating sequences of events then we have to ask: 'when do we equate one event with another similar one?' If, on the other hand, following Lewis [5], we see causation as counterfactual implication grounded in possible world semantics, we have the problem of the identity of counterparts in different possible worlds.

The issue of identity that we raise here goes a little deeper. When is one thing two things, and when are two things one thing? Our belief in the unity of the self (see, for example, Cupitt [3] for a critique of this belief) tends to draw us towards unity; our continual practice of dissection draws us towards plurality.

Paraconsistent formulations can be used to address these problems of identity. We have shown this by our paraconsistent formulation of the network stack. The same technique can be used more generally. (An ontologically problematic individual has a property if any of the individual's non-problematic instances has the property.)

6. Conclusion

We have shown that using multiple causal fields gives rise to inconsistency. We have indicated that a paraconsistent approach may allow us to live with these inconsistencies. This may motivate further work to create a framework for causal reasoning within a paraconsistent logic.

We have also raised the problem of identity, most clearly with reference to a message being passed down and up the OSI stack. We have proposed a general paraconsistent solution to this problem.

7. References

[1] Kenneth P. Birman & Thomas A. Joseph, Exploiting virtual synchrony in distributed systems. Proceedings of the Eleventh ACM Symposium on Operating Systems Principles, Austin, Texas, November 1987, pp 123-128.

[2] Jeremy J. Carroll & Andrei V. Borshchev, A Deterministic Model of Time for Distributed Systems, Proceedings of the Eighth IEEE Symposium on Parallel and Distributed Processing, New Orleans, Louisiana, October 1996, pp 593-598.

[3] Don Cupitt, The Time Being, SCM Press, London, 1992.

[4] David Hume, An enquiry concerning human understanding, Sect. VII. Of the idea of necessary connexion part II, 1748.

[5] D. Lewis, Causation, The Journal of Philosophy LXX, 1973, pp 556-567.

[6] J.L.Mackie, Causes and Conditions, American Philosophical Quarterly 2/4 70, (October 1965) pp 245-255 and pp 261-264.

[7] M.C.Menchini, A study of causation from the perspectives of philosophy and logic programming, MSc Thesis, Univ. of Manchester,

[8] Reiter, A logic for default reasoning, Artificial Intelligence 13, 1980, pp 81-132.

[9] Nicholas Rescher & Robert Brandom, The Logic of Inconsistency, Basil Blackwell, Oxford

[10] Shoham, Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence, MIT Press, Cambridge, Mass. 1988.

[11] Andrew S. Tanenbaum, Computer Networks, Third Edition, (International Edition), Prentice-Hall, 1996, pp 28-34.

[12] Zimmerman, Pretty Good Privacy: Source Code and Internals, Cambridge MA: M.I.T. Press, 1995.