FORMALISM HELPS IN DESCRIBING ACCIDENTS

Peter Ladkin, Universität Bielefeld, Germany

Karsten Loer, BAe Dependable Computing Systems Centre, University of York, UK

Abstract
We analyse the `probable cause' of the 1979 Chicago DC-10 accident using a minimal formalism, and find an omission. The omission is contained in the body of the report. This omission had consequences for the public discussion of this accident, which we show. We conclude that formalism helps in accident reporting by enabling simple consistency and omission checks. We then present a quick overview of our formal method, *Why-Because Analysis*, which provides the necessary mechanisms and rigor. We consider this to be the engineering of causal reasoning. As is now known from a quarter-century's experience with verification of digital systems, such reasoning engineering is both essential and non-trivial.

Accident reports in aviation present careful reasoned conclusions about causes and causal factors contributing to the accident, as well as providing *pro forma* details which may be useful in other contexts, say for statistical investigations of accident types. In order to establish the need for engineering this reasoning, we exhibit an example of the kind of problem that can arise through using the current informal methods. Second, we briefly describe our own "*reasoning engineering*" method for causal analysis, *Why-Because Analysis*, or WBA. WBA is similar in scope and rigor to verification methods familiar in computer science. It extends such methods in that it explicitly encodes a rigorous causal logic.

First, the statement of `probable cause' of the Chicago-O'Hare 1979 DC-10 accident (NTSB-AAR-79-17) :

> *The National Transportation Safety Board determines that the probable cause of this accident was the asymmetrical stall and the ensuing roll of the aircraft because of the uncommanded retraction of the left wing leading edge slats [...] and the loss of stall warning and slat disagreement systems resulting from maintenance-induced damage leading to the separation of the No. 1 engine and pylon assembly at a critical point during takeoff. The separation resulted from damage by improper maintenance procedures which led to failure of the pylon structure.*

We shall analyse this statement with the following semi-formal method. First, we list the crucial events, and denote each by a simple phrase. Second, we determine all relations between the events given by true assertions of the form: `*why* ..... *because* .... . (This second step has a formal semantics, as in Section *The Logical Semantics of Causal Explanation*, below, but it suffices for now to do it informally.)

A "Warm-Up" Example

Let us take first a simple hypothetical example to illustrate the method. An aircraft stalled, hit the ground, and was destroyed. That the aircraft was destroyed fits the definition of "accident" (significant damage to aircraft and/or loss of human life or severe injury). Let us apply the above steps.

Stage 1.
The crucial events are

- [0] the aircraft was destroyed
- [1] the aircraft hit the ground
- [2] the aircraft stalled

Stage 2.

Why *was the aircraft destroyed?*
Because *it hit the ground*

Using the notation above, this is `**Why** [0]? **Because** [1]', or, as we shall write it below,

[0] ~> [1]

That seems straightforward enough. Whether the event of the aircraft hitting the ground and the event of the aircraft being destroyed are the same "event" is something philosophers have argued about, but common sense use of the word suggests they're really two descriptions of the same event, so we shall go with this usage.

Why *did the aircraft hit the ground?*
Because *it stalled*

That is, [2] ~> [1]. Is this right? Did the aircraft hit the ground because it stalled? Well, actually no, maybe not. If the aircraft was at 100 feet, then probably yes. But suppose the aircraft was at 6000 feet. Aviation professionals will immediately spot a missing factor. The corrected relation is:

Why *did the aircraft hit the ground?*
Because *it stalled and did not recover in time*

We need to return to Stage 1 and modify.

Stage 1+.

The crucial events are

- [0] the aircraft was destroyed
- [1] the aircraft hit the ground
- [2] the aircraft stalled
- [3] Recovery was not effected in time

Stage 2+.

[1] ~> [0]
[2] /\ [3] ~> [1] (here, we use "/\" as formal notation for " *and* ")

Event [0] fulfils the FAA definition of accident. We have three events, one being the accident (remember that we took [0] and [1] both to designate the accident event), and the other two, [2] and [3], causally related to it, which we determined by asking `*why?....because*'. There are two causally determining events for [1], and both of these events had to happen for [1] inevitably to occur. Events [2] and [3] are thus jointly necessary but not individually sufficient for [1].

Let us summarise what we have done. We started to determine the "*why...because...*" relations amongst the facts/events, and found we had insufficient facts. We added a necessary causal factor (Stage 1+) and modified the derivation (Stage 2+). This revisiting is usual in formal methods, and could be named the `*inevitable intertwining*' or `*spiral*' (after Swartout and Balzer (SwBa82) , respectively Boehm (Boe86)). The final causal relation that we obtained in Stage 2+ holds as well for the airplane at 100 feet as for the airplane at 3000 feet. At 100 feet, noone expects a stall recovery to be effected in time, so one imagines it could be left unsaid. Our *why...because...* method seems to insist it be put in. Does it need to be said? Our view is, in an accident report, yes. A clear statement could lead to useful research into stall recovery in less than 100 feet. If it's not said, no-one will remark this possibility. (A parallel case might be the history of "pilot error" determinations, and the progress made in overall system design since one started asking *why* the pilots were making such errors, rather than accepting "pilot error" as a final reason beyond which no more need be asked.)

Prima facie, such a simple analysis could help us identify missing causal factors and ensure that our causal analysis is more correct. Let us now apply it to the statement of probable cause in the Chicago report. We shall see the same phenomenon, but with more details.

Analysing the Probable Cause of the Chicago DC-10 Accident

First, a list of events mentioned in the statement of probable cause:

- [1] The accident (= aircraft impacted ground and people on board died);
- [2] the roll of the aircraft;
- [3] the asymmetrical stall;
- [4] the uncommanded retraction of the leading edge slats;
- [5] loss of stall warning system;
- [6] loss of slat disagreement system;
- [7] separation of No. 1 engine and pylon assembly at critical point;

- [8] improper maintenance procedures.

Second, the apparent relationship between events as asserted in the `probable cause' appears to be a complex causal chain of the form

[8] ~> [7] ~> [5] /\ [6] /\ [4] ~> [3] ~>
[2] ~> [1] ~> [0]

So, the accident report considers a `probable cause' to be a causal chain. This is itself rather an enlargement of the concept of "cause". It singles out this causal chain as the most important interconnection of events. However, the stall warning system is an indication to the pilots of what was happening, as is also the slat disagreement system, and their loss ([5] and [6]) only affects at most pilots' behavior, and not directly the control systems of the aircraft. They certainly play no direct role in [3], [2] or [1]. Specifically, although ` *why [3]? because [4] /\ [5] /\ [6]*' is true, so is ` *why [3]? because [4]*'. Therefore one could conclude that [5] and [6] are superfluous in statement of *this* causal chain, since if it is a correct causal assertion, the following is also a causal chain leading to the accident:

[8]~>[7]~>[4]~>[3]~>[2]~>[1]~>[0]

However, during the discussion, the report says:

> *The simulator tests showed that, even with the loss of the number two and number four spoilers, sufficient lateral control was available from the ailerons and other spoilers to offset the asymmetric lift caused by left slat retraction at airspeeds above that at which the wing would stall. However the stall speed for the left wing increased to 159 KIAS.*

(KIAS denotes `*Knots Indicated Air Speed*', i.e., the figure displayed on the Air Speed Indicators in the cockpit.) The report is saying explicitly that [4] did not inevitably lead to [3]. The airplane remained controllable. That entails that [4] did not inevitably result in [3], i.e.,

[8] ~> [7] ~> [4]
[3] ~> [2] ~> [1] ~> [0]

which is no longer an unbroken chain. Three points are apparent:

- the causal chain leading to the accident is apparently then
- [3] ~> [2] ~> [1] ~> [0]
- in which [4], [5], [6] and their precursors do not appear;
- the statement of probable cause clearly intends to say that [5] and [6] were somehow involved;
- If by reasoning from the probable cause and the findings we can conclude both that [5] and [6] were not involved in the causal chain leading to the accident, and that they *were* indeed involved as causal factors, then this is inconsistent and there is a problem with the logic of the report.

The solution to this apparent inconsistency turns out to be, as before, that something is missing from the causal chain expressed in the `probable cause' statement. This omitted fact is, however, clearly described in the body of the report.

> *The evidence was conclusive that the aircraft was being flown in accordance with the carrier's prescribed engine failure procedures. [...] Since the wing and engine cannot be seen from the cockpit and the slat position indicating system was inoperative, there would have been no indication to the flight crew of the slat retraction and its subsequent performance penalty. Therefore, the first officer [the pilot flying] continued to comply with carrier procedures and maintained the commanded pitch attitude [...] which decelerated the aircraft towards V2, and at V2 + 6, 159 KIAS, the roll to the left began. [...] There would be little or no [impending-stall-indicating] buffet. [...] Since the roll to the left began at V2 + 6 and since the pilots were aware*

> *that V2 was well above the aircraft's stall speed, the probably did not suspect that the roll to the left indicated a stall. In fact, the roll probably confused them, especially since the stick-shaker [a stall warning] had not activated.*

This says clearly that because the flight crew were unaware of the slat retraction, they didn't know that the stall speed had increased, and they flew the airplane "*in accordance with procedures*" which dictated a speed slower than the new increased stall speed. It was thus inevitable that the airplane's left wing would stall. There was no indication to the pilots of this impending stall because the stall warning system was also inoperative. Had there been, one imagines that they would have reacted immediately (the indications are that they were excellent pilots, who the report says were flying exactly `by the book') and the airplane could have been controlled (the report has stated, above, that the airplane was controllable, derived from simulator tests).

Hence the report says that pilots' ignorance of the asymmetrical flap condition and impending stall allowed the stall of the left wing to take place. Thus there is an essential causal fact missing from the `probable cause' statement, namely:

[5] /\ [6] ~> [9] ~> [3]

where

[9] pilots continued to fly the airplane at below left-wing no-slat stall speed

Thus the causal chain should read

[8] ~> [7] ~> [5] /\ [6] /\ [4] ~>
[9] ~> [3] ~> [2] ~> [1] ~> [0]

which reinstates [5] and [6] to their intuitively proper places in the causal chain, and completes the chain as awaited.

So the logic of the report is faulty. The `probable cause' statement includes an incomplete causal chain. A simple semi-formal analysis of the report itself, namely just asking what the critical events were as expressed, and what the report says are their causal relationships, has exposed this incompleteness, and demonstrated the inconsistency in the report itself.

Consequences for Public Policy

Well, OK, an engineer might reply, maybe the report's reasoning doesn't satisfy the logical nit-pickers, but we can all figure this out from the report for ourselves, so why worry? The answer is that the statement of probable cause for significant accidents has consequences for public policy, and this policy may be mistaken in so far as the statement of probable cause is based on incorrect causal reasoning. Let us then take a quick look at what the public consequences were.

There was considerable public interest at the time concerning the engineering of the DC-10 because of the accident. McDonnell Douglas issued a report (McD79) in an attempt `*To Set The Record Straight*':

> *There is no point, as rule as old as Aristotle tells us, in debating a question that can be settled simply by examining the facts. [...]*
> *[The circumstances of the accident] gave rise to important - to urgent - questions. [Questions follow.]*
> *Naturally, properly, discussion of the DC-10 continued as long as such questions remained unanswered. And not all of them were answered quickly. […]*
> *The answers, when they emerged, were clear and conclusive. They proved that the DC-10 meets the tougest standards of aerospace technology. They proved, too, that the Chicago accident did not result from any deficiencies of aircraft design, and that steps taken shortly after the accident had eliminated any possibility of recurrence.*

In a section entitled **The Basic Questions**, they asked and answered:

- Why did a DC-10's pylon and engine separate from the wing at Chicago?

[Because of a very large crack in the horizontal flange of the pylon's aft bulkhead.]

- What was the origin of this crack? [Damaged by improper maintenance procedures, which were thereafter immediately `*banned by law*' as soon as discovered.]
- Have changes in the pylon's structural design been ordered? [No.]
- Is the pylon supported from the wing by a single quarter-inch bolt? [No.]
- Why were DC-10s grounded? [Premptive prophylactic action because of a failure to detect such cracks on other airplanes at the time of the accident, and subsequent detection of such cracks.]
- Are the DC-10's hydraulics systems effective and safe? [Yes.]
- Is there a problem with the DC-10's wing slats? [No.]
- But weren't changes to the slats required after the accident? [No. But stall-warning system changes were. They `*provide additional backup in the system [...]. The DC-10 stall warning system's "redundancy" - duplication to provide back-up security - exceeds industry standards for transport aircraft.*']
- [Some questions about ` *two other fatal DC-10 accidents in 1979 after the Chicago crash*'.]

McDonnell Douglas clearly felt the need to clarify public perceptions of the accident by enumerating and commenting the facts. This is a laudable goal, which we support.

First, we can imagine that a clear, consistent, complete explanation to the public of what had gone wrong, a goal of the NTSB, McDonnell Douglas, and the airlines, could have followed directly and unambiguously from the NTSB report without the intervention of McDonnell Douglas, had the NTSB report conclusion been complete and had the report itself not been inconsistent.

Second, McDonnell Douglas's `Basic Questions' generally follow the `probable cause' statement of the NTSB report. As factor [9] was not included from the `probable cause' statement, so it does not appear in the `Basic Questions'. An answer is given, however, namely that the stall warning system's redundancy "*exceeds industry standards for transport aircraft.*". We can conclude

- that the stall warning system redundancy did not suffice, since the airplane remained flyable, but the pilots flew it `*by the book*' into a stall;
- that if the redundancy `*exceeds industry standards*', the industry standards do not suffice.

The NTSB in fact drew both these conclusions, even though they do not explicitly pertain to the `probable cause' statement. The report's `*Safety Recommendations*' (Class II, Priority Action A-79-99) recommended that

> *[...] if certification is based upon demonstrated controllability of the aircraft under condition of asymmetry, insure that asymmetric warning systems, stall warning systems, or other critical systems needed to provide the pilot with information essential to safe flight are completely redundant.*

(This is the clause of A-79-99 pertaining to the DC-10. The McDonnell Douglas report states that the DC-10 was the only wide-body cabin airliner to have demonstrated the ability to fly with asymmetrical slats, which it did during certification.)

Conclusions About the Use of Formalism

This accident and report come from 20 years ago. One may ask, is this still relevant? We would answer yes, because while there has been considerable advance in the engineering of reasoning in, for example, computer science, especially with regard to distributed and safety-critical digital systems, none of this science has yet made it into other engineering domains such as accident analysis; we needed to motivate its introduction.

Our simple formalisation has shown infelicities in the NTSB report of its conclusions concerning the Chicago crash. McDonnell Douglas felt the need for public clarification, and a clear statement of the facts. However, full information on one necessary causal factor was not provided in their clarification. This is consistent with the omission of this factor from the statement of probable cause in the NTSB report. We can imagine that public and professional discussion of the accident, an essential factor for public risk assessment as it is now conceived (NRC96), could have been aided by simple formalisation, which demonstrates this omission.

This is not the only example to demonstrate advantages of this simple formalism. In (LaLo98), it was shown using the same technique that two necessary causal factors, the position of an earth bank and the state of the runway surface, were omitted from the `Causes' statement of the report on the A320 accident in Warsaw in September 1993. Both of these are under control of the Polish authorities, yet recommendations to the authorities were only that the system of collecting and distributing meteorological information should be adapted to conform to ICAO Convention Annex 3 standards, and that the bank should be described in the AIP Poland (the official description of airports). One can thus observe from the formalisation that the recommendation *prima facie* does not conform precisely to all the necessary causal factors, and imagine that it would have helped the goals of accident analysis to have addressed this apparent disparity in the report itself.

A WB-Analysis of the report of the American Airlines Accident in Cali in December 1995, also in (LaLo98), has found causal omissions, namely that the pilot-controller interactions did in fact causally contribute to the accident, contrary to what is stated in the report, but consistent with the NTSB's recommendations to the FAA concerning the accident.

We conclude that formalisation helps. It enables us to check not only the events, but also the reasoning concerning those events and the derivation of the conclusions and recommendations in an accident report. Now on to the proposal for engineering the reasoning.

Why-Because Analysis

Why-Because Analysis (WBA) is a method we developed for the failure analysis of complex, *open*, *heterogeneous* systems. The adjective `*open*' means that the behavior of the system is highly affected by its environment. Aviation operations are significantly affected by the weather through which aircraft fly, for example, and landing risks are significantly affected by obstacles and their clearance on the approach and go-around paths. The adjective `*heterogeous*' means that the system has components of different types that are all supposed to work together: digital, physical, human and procedural components, and combinations of some or all of these. Modern aviation operations have all of these components and thus form a complex, open, heterogeneous system.

The First Step - the WB-Graph Method

The WB-Graph (or WBG) Method develops the WB-Graph of the failure scenario, as in the examples above (the causal chain is called a "*graph*" because it may be drawn as a mathematical graph, and we have tools to do that - see below). The WB-Graph is a complete statement of the causal relations between all the events and system states of significance for causally explaining the failure scenario. The WBG Method consists broadly speaking of the two steps we illustrated above:

- *list:* make a list of all the events and states of significance;
- *determine causal relations:* determine the causal relations between all these, exhaustively, using the semantical test for `*causal factor*' explained in Section *The Logical Semantics of Causal Explanation*.

The first step is what an investigator does `in the field', and is not new. The second, however, is unique to the WBG Method. It is a rigorous, objective assessment of the facts which produces a causal explanation which is objectively justifiable. The approach used builds on Lewis's seminal studies in the logic of causal explanation, as explained below.

The Second Step - Verification

But how does one know that one has not himherself made errors when using the WBG Method? The full WBA includes

- *formal proof:* a formal proof method which enables one rigorously to prove
  1. that the causal relations asserted by the WB-Graph are correct, and
  2. that sufficiently many factors have been identified to provide a sufficient causal explanation of each identified fact that is not a root cause.

What is the purpose of formal proofs? In system verification, experience shows that attempting a formal proof leads an experienced analyst very quickly to problems and mistakes in the analysis. In computer science contexts, attempting formal proof has exposed mistakes in widely-used telecommunications protocols, in the design of digital processors used in avionics, in clock-sychronisation algorithms for digital flight control systems, in the security of transactions in distributed databases, and in the design of cache-coherence protocols for multiprocessors. A formal proof in WBA provides an objective assessment that the analysis performed in the WBG phase is correct and relatively sufficient.

The formal-proof step of WBA can, and should, be omitted until analysts and investigators have familiarised themselves with and feel comfortable using the WBG Method. WBA specialists in formal proof can always be brought in to help if particularly sensitive investigations require the level of certainty attained by formal proof and local expertise is insufficient.

The WBA formal proof method is based upon a leading proof method used in the specification and verification of digital systems, the hierarchical proof method of Lamport (LamTLA). All formal proof methods are technically involved and mathematical, require some specialist training to use, are best left to technical documentation (LamTLA) (LaLo98), and therefore we forego more detailed comment on the WBA proof method here.

Finding all the Facts - the Method of Difference

Finally, how does one know that one has identified all the facts salient for an explanation? One cannot be completely sure by use of logic alone -- unfortunately, one cannot rule out using pure logic the (rather silly!) possibility that a fuel-tank explosion was caused by little green men playing inside it with matches. So some judgement has to be used (Artificial Intelligence specialists confronted with the same issue call such judgement a `*closed world assumption*'). Nevertheless, most investigators will know of cases in which something has been missed. For example, had the nearly-intact FMS not been found at the site of the Cali accident, significant explanatory facts would not be known, and the explanation of the accident would be unsatisfactory and incomplete.

In the Cali case, before the discovery of the FMS, one knew that there were things about the navigational behavior of the aircraft that required explaining but which the available facts were insufficient to explain. But how about the cases in which we don't know that we don't know? WBA promotes use of a version of Mill's *Method of Difference* (Mil43), which is roughly the following.

- *MD:* To find a causal explanation of a significant fact F, ask how the system behavior would have been different, had that fact not pertained; say behavior B. Compare behavior B with the actual behavior (which includes F). Ask: where is the first significant place they diverge from each other? A thorough description D of that place (event or state) contains a causal factor of F. Try to identify it (say G). Repeat now with G, and so on.

Use of *MD* requires a certain skill. Astute readers will have observed that the description of *MD* above is somewhat unspecific. True, but it is methodical, and we observe that many competent investigators use it implicitly anyway. Readers may prefer to consider it just an explicit formulation of good investigatory practice. We would claim that there is virtue in making such practice explicit; what is explicit is also methodical.

Irreducible Uncertainty - State Diagrams

What does an investigator do, when faced with a situation in which there is uncertainty about what happened that cannot be eliminated? In a situation in which one does not have enough facts to determine uniquely what happened? WBA provides

- *alternative-description:* a method for delineating and representing the alternative possibilities as a *Predicate-Action Diagram* or PAD (LamTLA).

A PAD does not provide certainty about what happened -- by assumption, we're dealing with a situation in which that is not to be had -- but it does provide a precise delineation of the problem area, and so enables prophylactic measures for future avoidance of the problem area to be designed as reasonably and as precisely as possible. Use of the PAD to describe areas of uncertainty is somewhat technical and is explained in (LaLo98).

Human Actions -The PARDIA Classification

To analyse human operator behavior, WBA uses an information processing classification (Nor88). which we call PARDIA (LaLo98, Ch. 7, The PARDIA Classification). PARDIA stands for the sequence of stages in effective (or ineffective) situational response:

Perception -- Attention -- Reasoning -- Decision -- Intention -- Action

As with any other component of the heterogeneous system, humans receive input from other components (normally through their senses) and, as `output', influence other components of the system through action. The PARDIA classification classifies failure in the human component into one of the above components of the PARDIA sequence. Why are there six components to PARDIA (others use four, for example)? We have found that improvements to the system that may be suggested to prevent a recurrence vary significantly according to the category of error. We have found that to each of the six categories, particular different sorts of system improvement can be recommended. These are listed in (LaLo98, Ch. 7, The PARDIA Classification). We call PARDIA a *classification* rather than a *model* because we don't claim that humans behave *as if* they were PARDIA automata -- they may or they may not -- but we do claim that the classification helps significantly in determining which sort of prophylactic measures for avoiding repeat incidents would be successful.

Formal Specification of Procedures and Regulations

The final feature of WBA is that the relevant procedures for human operators, and the regulations which govern them, may be formally specified in exactly the same manner as the behavior of the physical or digital components of the system. Again, this is not new -- work in the 1980's in Artifical Intelligence showed that significant portions of legal statutes could be satisfactorily represented as a logic program. However, we are the first to apply such techniques to aviation procedures and regulations. The procedures and regulations, as well as the PARDIA classification, are specified all in the same formal language that the WBA proof system uses, thereby obtaining the benefits of using a `*wide-spectrum*' language for this task.

Tools

WBA is supported by a number of freeware computer and design tools. The most important are:

- *WB-Script* is a language for describing a WB-Graph in written form.

- *wb2dot* is a tool, available for direct use over the WWW, that takes input in *WB-Script* and automatically draws the WB-Graph in *Postscript*. *wb2dot* invokes a parser for *WB-Script* which produces input for the *dot* graph-drawing tool from AT&T Research, passes this input to *dot*, and then files the *Postscript* output for downloading by the *wb2dot*-user. *wb2dot* was developed in the *PERL* multiplatform programming language.
- Various editing and checking tools based on freeware document preparation tools (GNU Emacs pf-mode, LaTeX style-file *pf.sty*, DATR-check) are available.

Summary

Use of WBA, then, involves first the WB-Graph Method -- building the WB-Graph from the list of facts. To help the investigation assemble the required facts, WBA promotes the use of *MD*, a version of Mill's Method of Difference. To handle resulting irreducible uncertainties, WBA proposes a method using PADs. To classify human behavior, WBA uses PARDIA. Finally, an attempt at formal proof that the explanation pictured in the WB-Graph is correct, based on a description of the incident in a wide-spectrum formal language, will expose hitherto unnoticed hiatuses in the explanation, and successful proof finally provides an objective criterion for determining the correctness of the causal explanation proposed. To our knowledge, WBA is the only failure-analysis method which provides all these features. Tools are available to support various aspects of WBA.

Analyses Performed To Date

Seven WBAs have been performed to date and one is in progress. Five of these analyses used the WB-Graph Method alone. Three include a formal proof.

The Logical Semantics of Causal Explanation

The WB-Graph Method is based on a formal semantics for causality introduced by the philosophical logician David Lewis of Princeton University (Lew73), (Lew86ii). Roughly speaking, the semantics of Lewis for the assertion that *A is a causal factor of B*, in which *A*, respectively *B*, is either an event or state, is that *in the nearest possible world in which A did not happen, neither did B*. This relies on a notion from formal semantics of `possible world', best illustrated by example. Suppose my office door is open. But it *could have been* shut. A semanticist can now say: in another possible world, it *is* shut. A possible world is a way of talking about things that could happen, but didn't. But what about `near' possible worlds? The `nearest' possible world in which my door is shut is one in which my door is shut, air currents around it behave appropriately, sound through it is muffled as it should be, but broadly speaking everything else remains the same. A further-away world would be one in which someone else who is not me is sitting here typing, and an even further-away world is one in which this whole environment is situated in Ghana rather than Germany. Now, suppose my door shuts. What caused it to shut? I was pushing it shut. The air was still, there was no draft, the only thing moving was the door and it was moving because I was pushing it shut. Intuitively, my actions caused the door to shut. How do I know this from the formal semantics? In the nearest possible world in which I didn't push the door, did the door shut? We have already supposed that nothing else was moving, no air drafts, no other person in the vicinity, so in the nearest world these would also be the case. It could be that all the molecules in the door moved the same way at the same time, so the door spontaneously shut - but this situation is so highly improbable as to be almost unthinkable, so could it be really the *nearest* such world? No. In the nearest world, everything behaved the same way, except that I didn't push the door. So it didn't shut. So according to my formal semantics, my action caused the door to shut. This formal semantical test is particularly important in circumstances in which many causal factors conjoin to make something happen, which is by far the most usual case. The simple semantics asks a question of two events, or states, at a time, and by asking the question systematically of all pairs, pair by pair, a complex WB-graph may be systematically built.

References

(Boe86): B. W. Boehm *A Spiral Model of Software Development and Enhancement*, ACM SIGSOFT Software Engineering Notes 11(4):14-24, August 1986.

(FiBi92): John H. Fielder and Douglas Birch, eds., *The DC-10 Case*, State University of New York Press, 1992.

(LaLo98): Peter B. Ladkin and Karsten Loer, *Why-Because Analysis: Formal Reasoning About Incidents*, Document RVS-Bk-98-01, RVS Group, Faculty of Technology, University of Bielefeld, 1998. Available from `www.rvs.uni-bielefeld.de`

(LamTLA): Leslie Lamport,
The TLA Home Page, Available at
`www.research.digital.com/SRC/tla/`

(Lew73): David Lewis, *Causation*, Journal of Philosophy 70:556-567, 1973. Also in (Lew86).

(Lew86): David Lewis, *Philosophical Papers, Volume II*, Oxford University Press, 1986. Also in (Lew86).

(Lew86ii): David Lewis, *Causal Explanation*, in (Lew86).

(Mil43): John Stuart Mill, *A System of Logic*, London: Longmans, 1843, 8th edn., 1873. Quoted in (Ste97).

(McD79): McDonnell Douglas Corporation *The DC-10: A Special Report*, McDonnell Douglas Corp., 1979. Also in (FiBi92).

(Nor88): D. Norman *The Psychology of Everyday Things*, New York: Basic Books, 1988.

(NRC96): National Research Council, Committee on Risk Characterization, Paul C. Stern and Harvey V. Fineberg, eds., *Understanding Risk: Informing Decisions in a Democratic Society*, Washington, D.C.:National Academy Press, 1996.

(NTSB-AAR-79-17): National Transportation Safety Board, *Aircraft Accident Report, American Airlines, Inc. DC-10-10, N110AA, Chicago-O'Hare International Airport, Chicago, Illinois, May 25, 1979.*, Report NTSB-AAR-79-17, NTSB, Washington, DC, 1979. Also in (FiBi92).

(Ste97): Helen Steward, *The Ontology of Mind*, Oxford: Clarendon Press, 1997.

(SwBa82): W. Swartout and R. Balzer *The Inevitable Intertwining of Specification and Implementation*, Communications of the ACM 25(7):438-440, July 1982.