Part V

Less-Formal WBA of Important Incidents

Chapter 23

The Cali Accident

(with Thorsten Gerdsmeier and Karsten Loer)

This chapter is dedicated to the memory of Paris Kanellakis, who died with his family in this accident

We base WBA on the Lewis semantics for causality because, we have argued, it is precise and captures the intuitive meaning of 'cause' well enough for formal failure analysis. We justify this view in this and the next chapter by analysing two accident reports: those of the American Airlines B757 controlled flight into terrain (CFIT) accident in Buga, near Cali, Colombia in late 1995, and the Lufthansa landing overrun accident in Warsaw, Poland, in 1993. Rather than using the full WBA to establish the facts, we use precisely those facts cited in the body and conclusions of the reports, and seek only to establish the causal relations between them using the Lewis semantics. The results broadly agree with those in the reports, as we would expect, but use of the Lewis semantics highlights what we claim to be reasoning mistakes in the conclusions of both reports.

We performed these analyses before basing the full WBA on Lewis's counterfactual-conditional semantics, to assure ourselves that the semantics would be successful in this application. In our judgement, it succeeds, based on these two and other examples, and we took this as evidence of the suitability of the Lewis semantics as a basis for the full WBA method. We invite the reader to judge for hish-erself. We have included some motivating remarks that harp on topics already discussed and decided earlier. We find this nevertheless appropriate, because these examples contain evidence also for the decisions concerning the form of WBA. and pointing out that evidence helps to substantiate the view we take. Also, there are somewhat more complex occurrences of certain features, such as the event/process distinction, than occurred in the relatively simple example of NW052, and these require a more extended discussion.

23.1 The Accident

The accident took place on approach to Cali, Colombia, on 20 Dec 1995. The accident aircraft, an American Airlines Boeing 757-223, registration N651AA, hit mountainous terrain while attempting to perform a Ground Proximity Warning System (GPWS) escape manoeuvre, about 10 miles east of where it was supposed to be on the instrument arrival path to Cali Runway 19. (The GPWS warns pilots of imminent terrain proximity, and an escape manoeuvre is a procedure which pilots use to maximise their chances of avoiding the terrain.) Approaching from the north, the crew had been expecting to use Runway 1, the same asphalt but the reciprocal direction, which would require flying past the airport and turning back, the usual procedure. They were offered, and accepted, a ‘straight-in’ arrival and approach to Rwy 19, giving them less time and therefore requiring an expedited descent. The crew were not familiar with the ROZO One arrival they were given, became confused over the clearance, and spent time trying to program the Flight Management Computer (FMC) to fly the clearance they thought they had been given. A confusion over two navigation beacons in the area with the same identifier and frequency led to the aircraft turning left away from the arrival path, a departure not noticed by the crew for 90 seconds. When they noticed, they chose to fly ‘inbound heading’, that is, parallel to their cleared path. However, they had not arrested the descent and were in mountainous terrain. Continued descent took them into a mountain, and the GPWS (Ground Proximity Warning System) sounded. The escape manoeuvre was executed imprecisely, with the speedbrakes left out, as the aircraft flew to impact. The US National Transportation Safety Board believes that had the manoeuvre been executed precisely, the aircraft could possibly have cleared the terrain.

The aircraft should never have been so far off course, so low. The accident has been of great interest to aviation human factors experts. It was the first fatal accident to a B757 in 13 years of exemplary service.

We analyse the Cali accident sequence, using the system states and events noted in the accident report [Aer96]. Our causal analysis compares interestingly with the statements of probable cause and contributing factors in the report.

23.2 The Cali Report

The report concludes (p57):

```
3.2  Probable Cause

Aeronautica Civil determines that the probable causes of this accident
were:

1. The flightcrew’s failure to adequately plan and execute the
```

```
   approach to runway 19 at SKCL and their inadequate use of automation.

2. Failure of the flightcrew to discontinue the approach into Cali,
   despite numerous cues alerting them of the inadvisability of
   continuing the approach.

3. The lack of situational awareness of the flightcrew regarding
   vertical navigation, proximity to terrain, and the relative location
   of critical radio aids.

4. Failure of the flightcrew to revert to basic radio navigation at
   the time when the FMS-assisted navigation became confusing and
   demanded an excessive workload in a critieal phase of the flight.

3.3         Contributing Factors

Contributing to the cause of the accident were:

1. The flightcrew's ongoing efforts to expedite their approach and
   landing in order to avoid potential delays.

2. The flightcrew's execution of the GPWS escape maneuver while the
   speedbrakes remained deployed.

3. FMS logic that dropped all intermediate fixes from the display(s)
   in the event of execution of a direct routing.

4. FMS-generated navigational information that used a different naming
   convention from that published in navigational charts.
```

It is interesting to note that the probable causes are all stated as *failures* or a *lack*, that is, an absence of some (needed) action or competence. These are descriptions of persisting *state*. However, the accident is an *event*. Depending on what one counts as state and what as event, a causal sequence cannot normally contain one event alone.

An event is normally explained by the values of system state variables along with certain prior events. We may therefore suspect that the statement of probable cause in the report is logically inadequate because (at the least) incomplete. This suspicion may be substantiated by observing that all four 'probable causes' would have been true even if the aircraft had successfully executed the GPWS escape manoeuvre and landed safely later at Cali. Or had the faulty left turn away from the cleared airspace not been executed. A set of probable causes that allow the possibility that the accident would not occur is necessarily incomplete as a causal explanation, as may be seen by considering the Difference Condition and its discussion in Chapter 15.2.

In contrast to the four probable causes and four contributing factors of the report, the WB-graph contains 55 causally-relevant events and states mentioned in the report. The statement of probable causes and contributing factors is not intended to represent all causally-relevant events and states. However, we know of no generally-accepted logic-based methodology for discriminating 'important' causally-necessary factors from 'less important' causally-necessary factors, and do not propose any ourselves (we note that Lewis has a similar aversion, for what we take to be similar reasons [Lew73a]). Each factor shares the logical property that, had it not occurred, the accident would not have happened.

23.3 Linguistic Analysis of Pilot/ATC Communications

(with Dafydd Gibbon)

There are interesting linguistic aspects of the conversation between Air Traffic Control (ATC) and American Airlines 965 (AA965) shortly before the crash of AA965 near Buga, Columbia, on 20 December 1995. These comments derived from the NTSB Press Release of December 28th, 1995 [Boa95], and were originally published in [GL96].

The two relevant navigation plates are the *Cali VOR/DME/NDB Rwy 19 Instrument Approach Procedure* and the *Cali Rozo One Arrival Procedure* [Aer96, Appendix C].

We are not asserting here that the linguistic features we emphasise are causal to this accident. Whether they are will be determined by the official report of the accident investigation committee, and it would be inappropriate for us here to prejudge this determination. It is nevertheless clear that such features may be causally relevant.

First, there is *lexical ambiguity* in the ATC-procedural language: *ROZO* is the name of a *fix*[1] and also part of the name of an *arrival*[2] It's well-known that lexical ambiguities easily engender confusing conversations. It surprises us that, in a safety-critical situation such as an instrument approach procedure, names were not chosen to be unambiguous. Syntactically, '*ROZO*' is the fix and '*ROZO*

[1] *FIX- A geographical position determined by visual reference to the surface, by reference to one or more radio NAVAIDs, by celestial plotting, or by another navigational device* [U.Sb].

[2] *ARRIVAL ROUTES (ICAO): Route identified on an instrument approach procedure by which aircraft may proceed from the enroute phase of flight to the initial approach fix.* Jeppesen Chart Glossary, p5 (Aug 12-88).

One' is the arrival. Lexical ambiguity occurs because the word '*ROZO*' is used in two different lexical constructions[3].

Second, naming an *arrival* after its *endpoint*, rather than its *start*, may be semantically dissonant for US pilots. Semantic dissonance can easily lead to confusion with respect to temporal as well as spatial location (see, for example, [Gib]). In the USA, arrivals are generally named after starting fixes rather than near-endpoints, in contrast with naming of the *Rozo One Arrival*. However, we note that the opposite is true with road naming conventions in a variety of cultures (e.g., we expect the *Monterey Road* to lead in the direction of Monterey). The *Tulua* radio beacon is the entry point for the Rozo One arrival, some 22 miles before the initial approach fix for the Rwy 19 approach, itself 9 miles before the *Rozo* beacon, which is actually the *final approach fix*, 2.6 miles before the runway threshold.

Consequences of both these features are starkly present in the partial transcript provided by the Columbian authorities. Here is the report verbatim from the original clearance to the end of reported conversation. Our comments follow.

```
Approach replied, "Roger 965 is cleared to the VOR DME approach
runway one niner, ROZO Number One arrival, report Tulua VOR"

The flightcrew readback was, "Cleared the VOR DME one niner
ROZO one arrival, we'll report the VOR, thank you Sir"

Cali approach immediately clarified with, "Report Tulua", and
the flightcrew immediately acknowledged, "Report Tulua"

The flightcrew referred to the cockpit chart package (approach
publications) after ATC instructions to "Report Tulua"

Flightcrew discussion took place about the navigational aids to
be used in the ROZO 1 Arrival, specifically their position
relative to Tulua

About 30 seconds later the flightcrew requested, "Can American
Airlines 965 go direct to ROZO and then do the ROZO arrival
sir?"

Several radio transmissions then took place: Approach replied,
"affirmative direct ROZO one and then runway one niner, the
winds calm".  The flightcrew replied, "all right, ROZO, the
ROZO 1 to 19, thank you, American 965".  And the controller
```

[3] This confusion was also noted by Wally Roberts.

```
stated, "Affirmative, report Tulua and twenty one miles, 5000
feet".  The flightcrew acknowledged, "OK report Tulua, twenty
one miles at 5000 feet, American 965"
```

AA965's requested clearance (after the phrase '*About 30 seconds later*') demonstrates a linguistic expectation. Going '*direct to ROZO and then do the ROZO arrival*' only makes sense if one expects that fix *ROZO* is the beginning of the *ROZO* arrival. Which it ain't. That's point two at work.

Then, the response from ATC:

```
[..] Approach replied,
"affirmative direct ROZO one and then runway one niner, the winds calm"
```

First, the ATC response '*affirmative*' is incorrect. AA965 cannot fly what they requested. A *clearance*[4] is both a commitment by ATC to keep the airspace clear for the space and time slots indicated and a mandatory (but revisible) routing instruction. Any clearance must be a possible routing. Which AA965's request is not.

Second, notice the following words spoken by ATC: '*direct ROZO one*'. (This is a complete phrase, as indicated by the following conjunction, '*and*'.) It is a syntactically incorrect phrase in any grammar of pilot-controller speech. The word '*direct*' must be followed by the name of a fix. If '*ROZO*' is interpreted as the fix, the word '*one*' is then superfluous. The entire phrase '*ROZO one*' denotes an arrival, and an arrival name cannot correctly follow the word '*direct*'. The lexical ambiguity enabled production of this syntactically incorrect and semantically confused phrase.

Then follows a precise repeat from AA965 of their unfulfillable request, and a further '*affirmative*' from ATC.

A cognitive psychologist might note that, in order to make decision *D1: how to fly the arrival/approach*, AA965 had to make a decision *D2: which methods to*

[4] *AIR TRAFFIC CLEARANCE- An authorization by air traffic control, for the purpose of preventing collision between known aircraft, for an aircraft to proceed under specified traffic conditions within controlled airspace. The pilot-in-command of an aircraft may not deviate from the provisions of a visual flight rules (VFR) or instrument flight rules (IFR) air traffic clearance unless an amended clearance has been obtained. Additionally, the pilot may request a different clearance from that which has been issued by air traffic control (ATC) if information available to the pilot makes another course of action more practicable or if aircraft equipment limitations or company procedures forbid compliance with the clearance issued. Pilots may also request clarification or amendment, as appropriate, any time a clearance is not fully understood, or considered unacceptable because of safety of flight. Controllers should, in such instances and to the extent of operational practicality and safety, honor the pilot's request. FAR 91.3(a) states: "The pilot-in-command of an aircraft is the final authority as to, [sic] the operation of that aircraft." THE PILOT IS RESPONSIBLE TO REQUEST [sic] AN AMENDED CLEARANCE if ATC issues a clearance that would cause a pilot to deviate from a rule or regulation, or in the pilot's opinion, would place the aircraft in jeopardy* [U.Sb].

use to make D1. There are two main methods: *M1: conversation with ATC*, and *M2: looking at charts (including electronic charts).* By disavailing themselves of *M2* (consulting the charts), AA965 had only one method, *M1* (conversation with ATC) and no means of cross-check (under the assumption that these were the two main methods). This increases the chances of failure at *D1*. And, as we have seen, the *D1* decision failed because of the semantic failure of *M1*, the ATC-AA965 conversation.

23.4 The Textual Version of the Cali WB-Graph

We use an ontology of (partial system) states and events as described in Section 14.1. The sequence of events and states used in the graph are precisely those mentioned in the Cali accident report, with one exception. As discussed in Section 23.3 the cockpit voice recorder transcript shows that the crew asked for confirmation of a clearance that it was impossible to fly. The controller said '*Affirmative*', thus (falsely) confirming a clearance he knew to be confused and impossible to fly. The report mentions that the controller felt he was not able to explain to the pilots that they were confused. This was attributed to 'cultural differences'.

The NTSB recommendations [Boa96] also mention fluency training in Aviation English for non-native speakers.

As argued in [GL96], we take the semantics of ATC/pilot English literally. Had the controller responded with a '*Negative*', as required by the semantics, we believe that in the normal course of things the ensuing clarification would have enlightened the crew to their situational awareness problem and thus CFIT would have been avoided. The Lewis semantics thus requires us to consider this a causal factor. The repeated affirmation was simply a mistake. We explain this event by the factors 'textitcultural differences' and '*lack of fluency*' as in the report.

We remind the reader that state predicates are qualified with *since* (to denote an event after which these predicates remain true) and/or *until* (to denote an event before which the predicate has remained true). Two nodes are classified as processes, namely (1.2) and (1.2.1.3.1).

We constructed the textual form of the WB-graph as in Figures 23.4, 23.4, and 23.4. We annotate nodes representing pilot-behavioral failures with their classification according to the PARDIA scheme in Chapter 18. See [Mel97, Sea83] for further information and discussion of the relevant foundations for cognition and action.

```
[1]          /* AC impacts mountain */
   /\<-.1>  /* GPWS manoeuvre failed: since [1.1.1] */
   /\{-.2} /* AC in mountainous terrain: since [1.2.1.2] */

  <1.1> /\[-.1] /* GPWS manoeuvre initiated  */
        /\<-.2> /* AC did not exhibit optimal climb performance */
        /\[-.3] /* AC very close to mountains @ [1.1.1] */

    [1.1.1] [-.1] /* GPWS warning sounds */
      [1.1.1.1] [-.1] /* AC dangerously close to terrain */
        [1.1.1.1.1] /\[1.2.1.1]
                    /\[1.2.1.3]
                    /\<1.2.2.2>

    <1.1.2> /\<-.1> /* AC speedbrakes are extended: since [1.2.2.1.2] */
            /\<-.2> /* AC performs non-optimal pitch manoeuvre */

      <1.1.2.1> /\<-.1> /* CRW didn't retract speedbrakes according to procedure
                          //#ACTION# */
                /\<1.2.2.2.1>
        <1.1.2.1.1> <-.1> /* CRW unaware of extended speedbrakes
                           // #AWARENESS# */
          <1.1.2.1.1.1> <-.1> /* CD displays speedbrakes-extended  */

      <1.1.2.2> <-.1> /* PF doesn't hold optimal steady pitch attitude
                      // #ACTION# */
    <1.1.3> /\<1.2.1>
            /\<1.2.2>

  {1.2} /\<-.1> /* AC on wrong course/position (2D-planar): since [1.2.1.3.1.1] */
        /\<-.2> /* AC flying too low for cleared airspace (3rdD): since [1.2.1.3.1.1] */
```

Figure 23.1: The Textual Form of the Cali WB Graph (Part 1)

```
<1.2.1> /\[-.1] /* CRW turned to "inbound heading" at [1.2.1.3]
                // #DECISION# // #REASONING# */
        /\<-.2> /* CRW without situational awareness: since [1.2.1.3.1.1]
                // #PERCEPTION# */
        /\[-.3] /* AC arrived at (false) Position B: end of left turn */

  <1.2.1.2> /\<-.1> /* CRW unfamiliar with ROZO One Arrival and Rwy 19 Approach */
            /\<-.2> /* CRW high workload */
            /\<-.3> /* CRW used procedural shortcuts */
            /\[-.4] /* CRW request for confirmation of false clearance
                       twice confirmed by ATC */
            /\[-.5] /* FMC erases intermediate waypoints @ [1.2.1.3.1.1] */

    <1.2.1.2.2> /\<-.1> /* CRW must expedite arrival */
                /\<-.2> /* lack of external visual reference */
                /\<1.2.1.2.1>

      <1.2.1.2.2.1> /\<-.1> /* lack of time for executing arrival procedure */
                    /\[1.2.2.1.1]

      <1.2.1.2.2.2> /\<-.1> /* arrival takes place at night */
                    /\<-.2> /* few lighted areas on ground */

    [1.2.1.2.4] [-.1] /* ATC misuse of Aviation English */
      [1.2.1.2.4.1] /\[-.1] /* discourse under cultural dependencies */
                    /\<-.2> /* ATC lack of fluency in English */
                    /\<-.3> /* ATC lack of knowledge of AC position */

        <1.2.1.2.4.1.2> <-.1> /* Colombian ATC Use-of-English
                                 training/certification */
        <1.2.1.2.4.1.3> <-.1> /* no ATC radar coverage */

    [1.2.1.2.5] /\<-.1> /* FMC design */
                /\[1.2.1.3.1.1]

  [1.2.1.3] /\{-.1} /* AC left turn from true course for 90 seconds:
                       since [1.2.1.3.1.1] */
            /\<-.2>   /* CRW didn't notice left turn:
                       since [1.2.1.3.1.1]; until [1.2.1.3] */

    {1.2.1.3.1} /\[-.1] /* PNF gives 'R' to FMC */
                /\<-.2> /* FMC-database uses 'R' to denote ROMEO */
                /\<-.3> /* CRW didn't realize <1.2.1.3.1.2>:
                           since [1.2.1.3.1.1]
                        // #PERCEPTION# */
                /\<-.4> /* PNF didn't correctly verify FMS-entry:
                           since [1.2.1.3.1.1]
                        // #ACTION# */
```

Figure 23.2: The Textual Form of the Cali WB Graph (Part 2)

```
     [1.2.1.3.1.1] /\<-.1> /* CRW believes 'R' denotes 'ROZO' in FMC
                           // #ATTENTION# */
                   /\[-.2] /* CRW decides to fly direct 'ROZO' */

       <1.2.1.3.1.1.1> /\<-.1> /* ID 'R' and FREQ for ROZO on the approach plate
                                  correspond with an FMC database entry */
                       /\<-.2> /* ID/FREQ combination usually suffice to identify
                                  uniquely an NDB within range */
         <1.2.1.3.1.1.1.1> <1.2.1.3.1.2.1> /* ARINC 424 Specification */

       <1.2.1.3.1.1.2> <1.2.1.2.1> /* CRW unfamiliar with ROZO One Arrival
                                      and Rwy 19 Approach */

     <1.2.1.3.1.2> /\<-.1> /* ARINC 424 Specification */
                   /\<-.2> /* Jeppesen FMC-database design */

     <1.2.1.3.1.3> /\<-.1> /* FMC-displayed ID and FREQ valid for ROZO */
                   /\<-.2> /* CRW didn't perceive FMC-displayed Lat/Long
                           // #ATTENTION# */
       <1.2.1.3.1.3.1> <-.1>  /* ROZO and ROMEO have same ID 'R' and FREQ */
         <1.2.1.3.1.3.1.1> <-.1> /* Colombian government decision */

       <1.2.1.3.1.3.2> /\<-.1> /* FMC display figures small */
                       /\<-.2> /* CRW not trained to check Lat/Long */
                       /\<1.2.1.2.2>

     <1.2.1.3.1.4> /\<1.2.1.3.1.3.1>
                   /\<1.2.1.3.1.3.2>

<1.2.2> /\[-.1] /* AC starts expedited descent from FL230 */
        /\<-.2> /* AC expedited-descent continuous: until [1.1.1] */

    [1.2.2.1]  [-.1] /* CRW decision to accept Rwy 19 Approach */

    <1.2.2.2> /\[-.1] /* CRW extends speedbrakes */
              /\<-.2> /* CRW failed to arrest descent: until [1.1.1]
                      // #ACTION# */
      <1.2.2.2.2> <1.2.1.2> /* CRW without situational awareness:
                               since [1.2.1.3.1.1] */
```

Figure 23.3: The Textual Form (Part 3)

```
Glossary:

AC Aircraft
ARINC ARINC, Inc.
ATC Air Traffic Control
CD Cockpit Display
Course Two-dimensional straight-line ground track
CRW Crew
FLxyz Flight Level xyz = Altitude at which altimeter reads
   xyz00ft @ barometric setting 29.92"=1013hPa
FMC Flight Management Computer
FREQ (Navaid) radio Frequency
GPWS Ground Proximity Warning System
Heading Magnetic compass direction along which course is flown
ID (Navaid) Identifier (sequence of symbols)
Jeppesen Jeppesen-Sanderson, Inc.
Lat/Long Latitude and Longitude Values
Navaid Navigation Aid (radio beacon)
NDB Non-Directional Beacon (a navaid)
PF Pilot Flying
PNF Pilot Not Flying
ROMEO NDB near Bogota
ROZO NDB near Cali
Rwy xy Runway with heading xy0 degrees magnetic (to nearest 10 degrees)
```

Figure 23.4: Glossary of Abbreviations Used in the WB-Graph

23.5 Finding and Resolving Discrepancies

Checking the construction, we noticed that certain causal factors were missing. Namely, we noticed that the node

1.1.1.1.1 (*AC dangerously close to mountain*)

had no causal forebears. We noticed this during the course of listing the *source nodes* in the WB-graph. Causal factors with no causal forebears are precisely those nodes which do not have any in-edges, in other words, source nodes (see Section 23.9). Intuitively the node [1.1.1.1.1] should have causal forebears, since the aircraft was in the three-dimensional position it was in (physically) because of persisting course (2D) and altitude (1D) states, which were in turn consequences of certain command actions. A persistent course state is a consequence of (i) a particular heading flown (ii) from a given position; a persistent descent state was commanded at a particular point. Hence we looked for these events/states.

Looking over the textual form again, the causal forebears of [1.1.1.1.1] were already included. These reasons are: (course) [1.2.1.3] that the aircraft was at Position B; from whence [1.2.1.1] the crew turned to "inbound heading"; while ⟨1.2.2.2⟩ continuing their descent. We modified the textual description to include these three reasons for [1.1.1.1].

We also realised that reasons for ⟨1.2.1.2.1.1.1⟩ *CRW believes 'R' denotes 'ROZO' in FMC database* were in the report, respectively the NTSB recommendations, but had not yet been included in the textual graph: namely ⟨1.2.1.3.1.1.1.1⟩ *ID 'R' and FREQ for ROZO on the approach plate correspond with an FMC database entry* and ⟨1.2.1.3.1.1.1.2⟩ *ID/FREQ combination usually suffice to identify uniquely an NDB within range*. ⟨1.2.1.3.1.1.1.2⟩ has a reason already in the textual graph, namely ⟨1.2.1.3.1.2.1⟩ *ARINC 424 Specification*

We had already drawn the WB-graph (below) so we simply added the links, even though two links cross existing links, without attempted to make the graph planar (since we had one crossing link to begin with).

This experience confirmed our supposition that the textual form with path-numbering and pretty-printing is much easier to construct and check thoroughly than the graphical form of the WB-graph, in particular to check the correctness of the '*Why...Because...*' assertions themselves in terms of the counterfactual semantics; but that properties of the graphical form single out certain kinds of mistakes, such as source nodes (which represent the 'original causes', as noted below) which should nevertheless have causal forebears.

We concluded that the textual and graphical forms are complementary, that they are both needed for checking, and that therefore a WB-Graph construction method should involve always constructing both. If the full WBA is used, of course, these discrepancies would not arise in the same way – these conditions only occur when analysing previously-written accident reports using the Lewis causal semantics.

23.6 Automated WB-Graph Construction and Checking

We are aware of the possibilities of error when generating a moderately complex WB-graph such as this by hand. One of the authors of the report [GLL97a] on which this chapter is based, Thorsten Gerdsmeier, programmed the graph in DATR, a pure inheritance language developed for phonological analysis in computational linguistics. A DATR *theory* (program) is a set of *nodes*, with defined attributes and values; *queries* (requests for values of attributes) are processed by evaluating the attributes. Attribute values may be aliased to an attribute of another node, and there are defaults for evaluation.

Each state/event in the WB-graph was written as a DATR node, with value being the description of the state/event. Attributes are the reasons (corresponding to the indented bulleted list by the node name in the textual form; and in the graph itself the arrows of which this node is head), and also the nodes for which this node is a reason (occurrences of the node name in a bulleted reason-list in the textual form; and in the graph the arrows of which this node is tail). The whole forms a simple DATR 'theory' [Ger97].

The DATR theory is thus written using only local information about each node: its value (the description), ancestors (immediate causal factors) and offspring (nodes of which it is an immediate causal factor). We have already mentioned the principle that all event nodes must have at least one causal factor which is also an event, although states may have factors which are all states, or a mixture of states and events (see many papers in [ST93]). Such metaphysical restrictions can easily be programmed into the DATR representation. A DATR interpreter was used to run the following simple checks:

- does every node have at least one causal factor which is an event?
- is every node classified as exclusively either an event or a state?

We found four event nodes whose causal factors were all states, one state node which was written mistakenly once as an event, and one event node which was mistakenly written as a state. This consistency condition has global consequences. If an event is once mistakenly written as a state, then all causal factors in its history need only be states; whereas in fact the event must have at least one event as factor, and that event must have at least one, and so forth. When the mistake is found, and the 'state' rewritten as an event, a consistency check must be made on the entire history to make sure that each collection of factors for an event contains at least one event. Thus an error in miswriting an event node as a state, or an event node which has only states as causal factors, requires a consistency review on the entire subgraph 'backwards' from this node. Such errors are therefore expensive.

The fact that all three authors of the reporat had overlooked these simple and obvious inconsistencies in the 'carefully checked' textual version, and the cost in time of correcting them, established firmly for us the value of using such automated help in generating the WB-graph. We recommend that DATR be used according to the method of [Ger97], or some alternative simple automated consistency-checking tool be used, when generating WB-graphs of comparable or larger size. An alternative method of assurance of correctness is, of course, to use the full WBA with formal proof, but we anticipate many users will not want to go 'the full distance' on all their analysis tasks, and DATR encoding provides a useful alternative means of performing simple checks and building the graph.

23.7 'Processes' – Event/State Ambiguity

The intuitive semantics of the division into events and states is that an event represents an action, a once-only state change, and a state represents a persisting condition. At the 'level of granularity' at which reasons are considered in accident reports such as Cali, it may sometimes be difficult to tell if a condition should be classified as an event or a state. Such a node should normally be classified as a process, although we introduced the idea of processes in Section 14.1 as containing components of both events and states, not as a catch-all classification in case of indecision. However, our experience suggests that indecision normally arises because a given node appears to have elements of both states and events, not because it is devoid of either, so the classification as processes is appropriate. (Classifying certain nodes as processes may have consequences for the application of the consistency condition in the last section.)

For example, consider event *[1]*, the accident event. Its causal factors are two states, thus superficially violating the consistency condition. The second condition, (1.2), *AC in mountainous terrain* is a state as described; but what in fact caused the impact is that the aircraft was in the position it was with the flight path that it had, and this flight path intersected with the mountain. Having a particular position at a particular time may be regarded as an event, since it is more-or-less instantaneous; but it is expressed logically as a state predicate - it is not an action. The AC flight path, which is an AC state predicate, along with the position-time event-state will ensure that, in the absence of other intervening events, other predictable position event-states will occur in the future. Some causal factors such as (1.2) thus represent imprecise features of the flight which at this level of granularity may be classified as an event or a state – and thus as a process. Looking back over this argument should reassure the reader that it is appropriate so to classify the node.

There are precisely two such nodes in the Cali WB-graph:

- (1.2) *AC in mountainous terrain: since [1.2.1.2]*

- (1.2.1.3.1) *AC left turn from true course for 90 seconds: since [1.2.1.3.1.1]*

It is the 'level of granularity' at which the reasons are expressed which engenders this event-state ambiguity and requires us to classify the nodes as processes, rather than any fundamental problem with the ontology or our method. We note in support of this claim that (1.2) is very closely related to *[1.1.3]*, *AC very close to mountains at [1.1.1]*. These position nodes are obviously not independent.

In the case of the Cali accident report, event-state ambiguity only occurs with position/flight path factors. One can extrapolate and suppose that this will happen with other accident explanations also. Thus we recommend that all position/flight path factors in accidents be examined to see whether they should be classified as processes, as pure events, or as pure states. We do not know if there are other such specific features of accident explanations which require resolution as processes. We recommend attempting to classify nodes first as events or states, and then only classifying the node as a process if neither identification is suitable, because classifying as a process renders the simple consistency check above inapplicable. It is best, we think, to retain as many simple consistency checks as possible.

A more detailed ontological analysis of flight path/position dynamics should obviate the need for processes in this example. Introduction of the relevant mathematics of dynamics, however, would in our opinion be 'overdoing it' at this level of granularity: one does not need to know the precise physics in order to know that being too close to the mountains was a causal factor. But we feel it would be preferable to bring the dynamical theory and the ontology we use into a more close relation with practice: we recommend avoiding classification of nodes as processes because

- the classification is intuitive: we give no guidance on when to classify factors as processes, other than to say that it may be expected with position/flight path parameters;
- introducing a process relaxes the applicable consistency check on that node. We believe that in formal methods in general, the more thorough consistency checking may be applied, the better; we thus dislike weakening consistency checks, on principle.

23.8 The Cali WB-Graph Layout

Some attempt was made to construct a planar graph by hand, since at that time we did not have the *wb2dot* tool available. We explain the layout procedures which we used. There are two crossings in this WB-graph. We concluded that at least one was necessary (by cases, trying to eliminate it), so felt that two did not greatly lessen legibility. We attempted to make the graph planar by using the algorithm below. The algorithm involves calculating the '*relative shape*' of certain tree-subgraphs, and '*laying them out*'. This procedure is not exact, because aesthetic, readability and size criteria come into play, and these criteria may well be in conflict. Such a conflict can only be resolved on a case-by-case basis by prioritising the criteria. We indicate the relative-shape and layout techniques we have found useful, with the understanding that they *can* be, but *must not* be, followed.

The question might arise: why not use one of the planar-graph algorithms already in the literature? We have found that WB-graphs have roughly the form of a tree-with-links. In the present case, the number of link nodes is roughly a quarter of the number of nodes, and the number of links roughly half the number of link nodes. In other words, the links are relatively independent, and the crucial events and states have relatively independent causes. We are handling a real example, our technique is relatively simple to grasp, and sufficed. Furthermore, we could use it 'by hand' while preserving many of the characteristics of the layout; and even with corrections we only had two crossings, which remain completely readable. So we didn't see the need to use a more mathematically precise algorithm.

The graph construction proceeds as follows.

- The textual form is constructed as a tree, as above;
- Links to other nodes occur as leaves, as above;
- The graph containing *just the links* (the dashed edges in the WB-graph below) is drawn, the nodes labelled with their path-numbers;
- nodes labelled with the *greatest common subsequences* of the node labels are added (recursively);
- cycles were noted, and for each cycle a list of links from nodes exterior to the cycle to nodes interior to the cycle (and vice versa) was made;
- an attempt was made to planarise the graph by 'inverting' cycles: writing the clockwise node-sequence as an anti-clockwise node sequence instead, whereupon some links from exterior to interior can become purely exterior links (and vice versa);
- '*leaving enough space*' to draw in the rest of the graph (which consists of pure trees), and drawing it in.

The '*leaving enough space*' technique consists in the following. The rest of the graph consists in a pure tree structure. Nodes in the link graph therefore have trees 'hanging off' them. The *relative shape* of each 'hanging' tree is calculated, and a decision is made (arbitrarily, that is, for aesthetic reasons) whether to 'hang' the tree off the exterior or interior of any cycle in which the mode partakes. The relative shapes of all tree structures in the interior of each cycle are *laid out* without overlapping, and the cycle is drawn outside this layout.

The *relative shape* calculation technique is as follows. Let one (space) unit be the length of a normal link between two nodes. A tree with n nodes of depth less than or equal $int(log2(n))$ may be drawn in a triangle of height $int(log2(n))$ and base $int(log2(n))+1$, where $int(n) \triangleq$ *the greatest integer less than or equal to* n. This is roughly the shape of the full binary tree with n nodes: it is our experience that any roughly 'bushy' tree with n nodes can be fitted in to such a shape without significantly affecting readability (not all nodes with a common parent will appear on the same level, so some links will need to be stretched). A 'linear' tree with n nodes is roughly the shape of a chain of length n, that is, a rectangle with width one node and length n nodes (the link to the parent in the link graph is included in the shape). Such a rectangle can of course be 'kinked'.

The '*layout*' algorithm consists in taking the '*chains*' and '*bushes*' and arranging them without overlapping as desired, kinking the chains as need be.

There are 59 nodes in the WB-graph for the Cali accident, with 14 'links' between 22 nodes (as may be easily seen from the textual form). The 'link graph' includes a few more nodes, but remains roughly half the size of the complete graph. The 'relative space' algorithm therefore is easy to use, since the trees to be 'hung' are relatively small. The final graph we drew has two crossings, one of them due to the two late corrections mentioned above. We did not think there was much point in trying to redraw the graph to see if we could eliminate the 'new' crossing. Using this algorithm for the Cali accident sequence, with corrections, yields the graph in Figure 23.5.

23.9 Source Nodes in the WB-Graph

In principle, *source* nodes (nodes with only outgoing edges, no incoming edges) represent reasons for the accident which have no further reasons lying behind them. They should thus represent the original reasons for the accident. Since the semantics of the WB-graph is that the node at the tail of an edge represents a necessary causal factor for the node at the head of the edge, source nodes represent necessary causal factors for the accident which themselves have no necessary causal factors mentioned in the report. Logically, therefore, the report regards these as contingencies, that is, events or states which need not have occurred, but whose conjunction was sufficient to ensure that the accident happened. These 'original causes' are as follows. (Notice that being an 'original cause' does not imply temporal priority - some original causes occur late in the accident event sequence.) The source nodes are displayed in Figure 23.6.

23.10 Critique

We find that the source nodes correspond pretty closely to what intuitively one could take as 'original causes' of the accident. This is not to say that the actions mentioned in this list were all unwarranted. For example, extending the speed-brakes and leaving them out was necessary to get down fast. But this *in combination with* the course change led to a fatal excursion out of protected airspace. One may observe that three of these source nodes, namely, the Colombian government decision on beacon ID/FREQ, the cultural dependencies in ATC/CRW discourse, and the Colombian ATC Use-of-English training/certification, were emphasised in the NTSB recommendations but not in the final report.

This list of source nodes could be used as follows. Procedures could be developed that avoid this fatal combination of circumstances. Exactly which procedures is a matter for expert judgement. For example, maybe airlines should not fly into Cali at night (expert judgement would not in fact be likely to draw this conclusion from this accident alone, given the other combination of factors). Some of these circumstances are already legislated against (being unfamiliar with the approach; accepting an approach which one has a lack of time adequately to execute, being unaware of extended speedbrakes). Avoidance is a matter for enhanced training, as is the non-optimal pitch profile flown in the GPWS manoeuvre. Enhanced ATC English and discourse training is also indicated. Pilot procedural modification is indicated: pilots should check Lat/Long. Technical modification is indicated: Enhanced GPWS, for example; maybe more perspicuous indication of Lat/Long on the FMC; maybe more perspicuous database notational standards; maybe modification of the ROZO ID/FREQ; maybe radar coverage in the Cali area. On the other hand, the NTSB pointed out that Cali is the only location they found worldwide in which the ID/FREQ combination does not suffice uniquely to identify an NDB within a radio reception area [Boa96], which would point towards modification of the ROZO ID/FREQ as being an appropriate response. Many of these issues are explicitly addressed in the report recommendations and the NTSB recommendations. This gives us confidence that our formal approach is consistent with the judgement of experts, while enhancing the ability to check for completeness and consistency of an accident explanation.

We are somewhat concerned that an intuitively important component of the crew's cognitive state, namely

```
<1.2.1.3.1.1.1>CRW believes 'R' denotes 'ROZO' in FMC
```

does not occur in the 'original cause' list. Our concern arises because, while one may inquire what encouraged them to hold this false belief (namely that ID/FREQ usually suffices for unique identification, and the ID/FREQ combination they chose corresponds with what is on the paper Approach plate), which does more-or-less completely explain the false belief, it does not vindicate the pilot behavior. They could have cross-checked better (the Lat/Long; being extra-aware whether the aircraft was turning away from course; gaining altitude while completing the cross-check). This case points out an important caveat.

It is important to realise that even if an event or state it has reasons in the WB-graph, not all the reasons may be included. Only a full WBA can guarantee a sufficient set of reasons for each internal (not-a-source) node. This report-analysis method based on the Lewis semantics alone and not the full WBA takes as nodes only those which have been identified informally by the experts. Although reasons for the crew's mistaken belief about 'R' denoting ROZO are given, some are missed out (as noted above). Procedurally, this mistaken belief could have been avoided by appropriate checking, and the checking that might be deemed appropriate might extend beyond the Lat/Long checking. In fact, the report recommends that procedure dictates that a go-around should have been performed in this situation. We did not perform a formal comparison with procedure in this particular application of the WB-graph method, as we did in Section 22. Such a comparison is in general necessary.

One may further observe that paying attention to the source nodes alone might cause one to miss the wood for the trees. For example, responding directly to the night flight or lack of lighted objects on the ground, one might try to prohibit night flights into Cali, or light up the surrounding countryside. But these measures would not help during extensive cloudy weather. The crucial component here is reduced visibility, interior node $\langle 1.2.1.2.2.2 \rangle$ in the WB-graph. A more careful method for evaluating causal components of the accident, then, would also look upstream from the source nodes to identify more general themes, such as lack of visual reference, that could be addressed by legislation or training or other means.

We may conclude that identifying source nodes is an important component of accident analysis using the WB-graph, but that techniques for comparison with procedure; and source nodes, while being the true ‘original causes’, may themselves describe circumstances too specific to dictate appropriate avoidance responses: one must check also further up the WB-graph for the most appropriate description of circumstances for which to formulate an avoidance response.

23.11 Discriminating the ‘Significant’ Events

The WB-graph method as presented here does not incorporate any mechanism to indicate the relative weight attached to events and states. However, in order reasonably to assess the accident, such weights must be given, as demonstrated clearly to us by Barry Strauch:

> *[....] the [WB-graph] methodology does not appear to give enough weight to how the crew's action in taking the controller's offer to land*

> *on 19 constrained their subsequent actions. [....] not all decisions are equal at the time they are made, [....] each decision alters the subsequent environment, but that while most alterations are relatively benign, some are not. In this accident, this particular decision altered the environment to what became the accident scenario.*
> [Str97]

Intuitively, this decision led directly to the crew's high workload, and also, because of their unfamiliarity with the arrival and approach, to their loss of situational awareness, communication confusion, and lack of attention to indicators of their situation: in short, to most subsequent causal factors.

Strauch notes that, for example, because of this decision, the workload was such that the crew failed to look at the Electronic Horizontal Situation Indicator (EHSI), which clearly and continually indicated the event-state (1.2.1.3.1), the continual left turn towards ROMEO, on a moving-map style display . The EHSI is one of two largish electronic displays in front of both pilots (the other displays physical flight parameters). Strauch notes that, because of this display,

> *[....] the interpretation of the effects of [the execution of [1.2.1.3.1.1]] should have required almost no cognitive effort. This is, in fact, one of the substantial advances of "glass cockpit" aircraft over older ones. As a result, regardless of the considerable effort required to verify R through the lat/long coordinates in the CDU, the EHSI presentation of the projected flight path displayed the turn.*
> [Str97]

Thus the high workload was such that even very easy cognitive tasks were significantly impaired, which is not indicated in state << 1.2.1.3.2 >>, the relevant *Awareness Failure*. The point of accident analysis is to determine what changes could be made in the future to avoid similar events and situations. The force of Strauch's point is that the crew were not just highly-loaded, but in 'cognitive overload'. Since they were in cognitive overload, the suggestion to modify training requirements to emphasise paying more careful attention to the EHSI, say, would not help avoid a repeat; probably this could not cognitively have been accomplished. However, determining such a situation definitively is beyond current expertise. The appropriate action is to emphasise decision-making methods that avoid the crew putting themselves in a situation of cognitive overload, and that get them out of such situations quickly if they feel themselves entering one. A basis for determining this difference in prophylactic action must be given by any complete accident analysis method. The WB-graph method thus requires a means of identifying such significant events as the acceptance of the ROZO One-Rwy 19 arrival and approach, as Strauch suggests. How may we do this?

We make a few informal observations. The decision to accept the ROZO One-Rwy 19 arrival and approach altered the goal of all subsequent actions, and thus in many cases those actions themselves. In the formal ontology used in the WB-graph method, a behavior is a sequence of exactly interleaved states and events: state-event-state-event-.... and so forth. Some of these behaviors fulfil the requirements (to land safely and normally on Rwy 19) and some of them (the accident sequence, for example) do not. At the point at which the ROZO One-Rwy 19 was accepted, the past behavior of the aircraft formed a definite, finite state-even-state-event-... sequence, which would be completed by one of a large number (possibly infinite, depending on how deep into the analysis one goes) of possible future behaviors. At this point, then, the state-event-state-event sequence looks like a sequence in the past and a tree in the future. (This is the semantics of, for example, the temporal logic CTL used in verification of concurrent algorithms, and a similar structure to the many-worlds interpretation of quantum mechanics.)

At the point of the ROZO One-Rwy 19 acceptance, the future behaviors satisfying the goal of the flight all consist in a safe landing on Rwy 1, and the future such behaviors after the acceptance all consist in a safe landing on Rwy 19. Not only are these sets of future behaviors disjoint (there is no behavior which belongs to both future trees), but they are radically disjoint - most of the events and states occurring along a future branch of the Rwy 1-tree would not occur along a future branch of the Rwy 19-tree. This radical disjointness property is precisely that which formally corresponds to ‘altering the environment’ of the flight, in Strauch’s words. The formal problem is then to find some logically and computationally sufficient means of assessing actions for the determination of ‘radical disjointness’ of their future trees. This is a question for further research.

23.12 A Comparison with the Cali Conclusions

We tabulate and compare the conclusions of the Cali report with the WB-graph analysis. The conclusions of the Cali report may be found below in Section 23.14.

Findings 2, 13-14, 16-18 are outside the scope of the WB-graph. They concern the general procedural environment in which aviation is conducted, whereas the WB-graph concerns itself only with the immediate actions and states in the time interval during which the accident sequence occurs. Finding 1, on the other hand, consists of the pro forma statement that the pilots were trained and properly qualified, conjoined with a statement that they suffered no behavioral or physiological impairment. The latter conjunct is in the domain of reference the WB-graph - it states that a condition pertained which, had it not pertained, could have helped explain the accident (and thus altered the form of the WB-graph). As far as the WB-graph is concerned, it is of the same form as the assertion that all aircraft systems, indeed the aircraft itself, worked as designed and intended. This assertion is a state predicate which remains true for the entire accident sequence, and clearly has causal consequences: had the systems malfunctioned somehow, the WB-graph would have looked different. However, when using the Lewis semantics for counterfactuals to evaluate the edges in the WB-graph, the 'nearest possible worlds in which ... is not true' are always those in which the systems functioned normally and the pilots suffered no impairment. We choose not to complicate the representation of the WB-graph by including these 'environmental assertions', but that should not be taken to imply that we do not consider them causally relevant.

Finding 6 is a consequence of the crew's regulatory-procedural environment. It is a general requirement on flight crew in the US, Western Europe and other ICAO countries that the report judges was not adhered to in the Cali incident. However, it is not directly causal, like the violation of other procedural requirements, and does not appear explicitly in the graph, while nevertheless being related to ⟨1.2.1.2.3⟩, that the crew used procedural shortcuts. It is, of course, important for explaining an accident that certain normative requirements were not adhered to, because the purpose of explaining an accident is to determine what may be changed in the future to prevent a repetition of similar incidents. If regulations were broken, that indicates that appropriate regulatory safeguards were already in place. We have suggested a method for identifying and including conflicts with normative requirements in [PL97], but don't apply it here, partly because the method is not yet fully developed and we feel that the Cali case is a more complex application which we prefer to address later.

The correspondence of the report's other findings, 3-5, 7-12 and 15, with items in the WB-graph is as follows:

Report Finding	WB-graph entry
3	[1.2.2.1.1]
4	⟨1.2.1.2.2.1⟩
5	⟨1.2.1.2.2.1.1⟩
7	⟨1.2.1.2.1.2⟩, ⟨1.2.1.3.1.1.1⟩, ⟨1.2.1.3.1.2.2⟩, ⟨1.2.1.3.1.3.1⟩
8	⟨1.2.1.3.1.3⟩, ⟨1.2.1.3.1.1.1⟩
9	[1.2.1.3.1.1], ⟨1.2.1.3.1⟩, ⟨1.2.1.3.1.4⟩, ⟨1.2.1.2.3⟩
10	⟨1.2.1.3.1⟩, [1.2.1.1]
11	⟨1.2.2.2⟩
12	⟨1.1.2.1.1.1⟩
15	[1.2.1.2.4]

We remark that Finding 15 and [1.2.1.2.4] correspond - because they are in contradiction! We noted above, however, that the causes of [1.2.1.2.4] are explicitly addressed by the NTSB Recommendations [Boa96]. We conclude that the report and the NTSB Recommendations do not concur on this event or its causal factors, and we chose to follow the NTSB view, as explained earlier and as argued by the second author (prior to the NTSB Recommendations) in [GL96].

With nearly sixty nodes, only sixteen of which correspond to the report's findings (of which there are only 10 pertinent to causally-relevant states or events in the domain of the graph), we conclude that the WB-graph yields a more thorough classification of the causally-relevant findings of the Cali accident investigation commission than the Findings Section (3.1) of the report.

23.13 Conclusions from the WB-Graph Construction

We have analysed the causal explanatory relations between the events and states listed in the Cali accident report [Aer96]. We identified 59 causally-relevant and -necessary factors, and constructed the WB ('explanatory') relation between them. We represented the result in textual, then graphical form. We found it easier to construct the WB-graph in this fashion.

We found that the list of 'source nodes', a conjunction of necessary and sufficient causes for the accident that were themselves regarded as contingent, is a fairly accurate indication of the causes, but should be used as guidance, and not uncritically, in formulating statements of cause and contributory factor. WBA does not yet include any method for weighing the relative importance of causal factors, so may not be used alone for distinguishing probable cause from contributory factor, or for assessing the comparative global significance of actions such as the decision to accept the ROZO One-Rwy 19 arrival and approach.

This application of WBA, constructing the WB-graph alone, is based on application of a rigorous logical criterion of explanation applied to the events and states identified by domain experts as crucial to the accident. The result is a data structure, the WB-graph, expressed in two forms, textual and graphical, each with their own analytic advantages. Both structures are manageable, as we have demonstrated on a real example. However, automated help, such as that provided by implementation in DATR, is highly recommended, both to avoid local errors and save the resources required to determine and correct their global consequences.

Furthermore, the graph represents 59 states and events noted by the Cali accident investigation commission and the NTSB as being causally-relevant. In contrast, the reports Findings section lists only 16 of these (roughly a quarter), corresponding to 10 explicit findings.

23.14 The Cali Report Conclusions

From [Aer96]:

```
3. 1  Findings
```

1. The pilots were trained and properly certified to conduct the flight. Neither was experiencing behavioral or physiological impairment at the time of the accident.

2. American Airlines provided training in flying in South America that provided flightcrews with adequate information regarding the hazards unique to operating there.

3. The AA965 flightcrew accepted the offer by the Cali approach controller to land on runway 19 at SKCL.

4. The flightcrew expressed concern about possible delays and accepted an offer to expedite their approach into Cali.

5. The flightcrew had insufficient time to prepare for the approach to runway 19 before beginning the approach.

6. The flightcrew failed to discontinue the approach despite their confusion regarding elements of the approach and numerous cues indicating the inadvisability of continuing the approach.

7. Numerous important differences existed between the display of identical navigation data on approach charts and on FMS-generated displays, despite the fact that the same supplier provided AA with the navigational data.

8. The AA965 flightcrew was not informed or aware of the fact that the "R" identifier that appeared on the approach (Rozo) did not correspond to the "R" identifier (Romeo) that they entered and executed as an FMS command.

9. One of the AA965 pilots selected a direct course to the Romeo NDB believing that it was the Rozo NDB, and upon executing the selection in the FMS permitted a turn of the airplane towards Romeo, without having verified that it was the correct selection and without having first obtained approval of the other pilot, contrary to AA's procedures.

10. The incorrect FMS entry led to the airplane departing the inbound course to Cali and turning it towards the City of Bogota. The subsequent turn to intercept the extended centerline of runway 19 led to the turn towards high terrain.

11. The descent was continuous from FL 230 until the crash.

12. Neither pilot recognized that the speedbrakes were extended during the GPWS escape maneuver, due to the lack of clues available to alert them about the extended condition.

13 Considering the remote, mountainous terrain, the search and rescue response was timely and effective.

14. Although five passengers initially survived, this is considered a non survivable accident due to the destruction of the cabin.

15. The Cali approach controller followed applieable ICAO and Colombian air traffic control rules and did not contribute to the cause of the accident.

16. The FAA did not conduct the oversight of AA flightcrews operating into South America according to the provisions of ICAO document 8335, parts 9.4 and 9.6.33.

17. AA training policies do not include provision for keeping pilots' flight training records, which indicate any details of pilot performance.

18. AA includes the GPWS escape maneuver under section 13 of the Flight Instrument Chapter of the Boeing 757 Flight Operations Manual and Boeing Commercial Airplane Group has placed the description of this maneuver in the Non Normal Procedures section of their Flight Operations Manual.

3.2 Probable Cause

Aeronautica Civil determines that the probable causes of this accident were:

1. The flightcrew's failure to adequately plan and execute the approach to runway 19 at SKCL and their inadequate use of automation.

2. Failure of the flightcrew to discontinue the approach into Cali, despite numerous cues alerting them of the inadvisability of continuing the approach.

3. The lack of situational awareness of the flightcrew regarding vertical navigation, proximity to terrain, and the relative location of critical radio aids.

4. Failure of the flightcrew to revert to basic radio navigation at the time when the FMS-assisted navigation became confusing and demanded an excessive workload in a critieal phase of the flight.

3.3 Contributing Factors

Contributing to the cause of the accident were:

1. The flightcrew's ongoing efforts to expedite their approach and landing in order to avoid potential delays.

2. The flightcrew's execution of the GPWS escape maneuver while the speedbrakes remained deployed.

3. FMS logic that dropped all intermediate fixes from the display(s)

in the event of execution of a direct routing.

4. FMS-generated navigational information that used a different naming convention from that published in navigational charts.

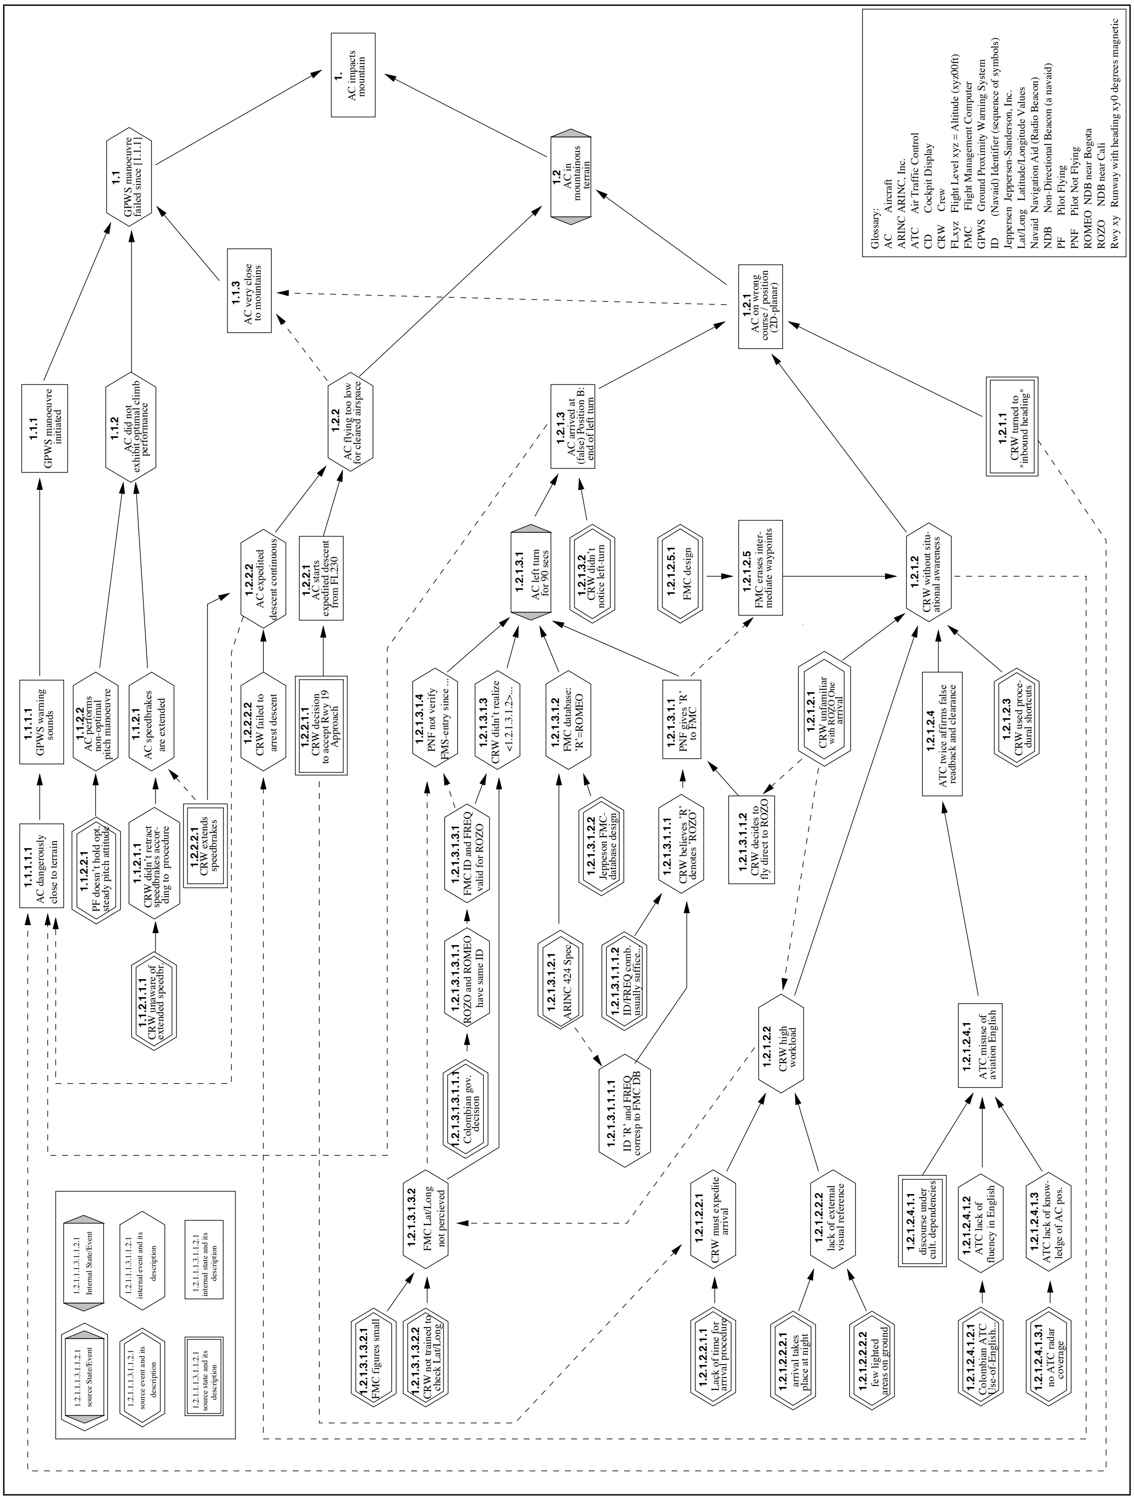

Figure 23.5: The Cali WB Graph

```
<1.1.2.2.1> /* PF doesn't hold optimal steady pitch attitude in
               GPWS manoeuvre
               // #ACTION# */
<1.1.2.1.1.1> /* CRW unaware of extended speedbrakes in GPWS manoeuvre */
[1.2.2.2.1] /* CRW extends speedbrakes for descent */
[1.2.2.1.1] /* CRW decision to accept Rwy 19 Approach */
<1.2.1.3.2> /* CRW didn't notice left turn caused by FMC:
               since [1.2.1.3.1.1]; until [1.2.1.3] */
<1.2.1.3.1.2.1> /* ARINC 424 Specification */
<1.2.1.3.1.2.2> /* Jeppesen FMC-database design */
<1.2.1.3.1.3.2.1> /* FMC display figures small */
<1.2.1.3.1.3.2.2> /* CRW not trained to check Lat/Long on FMC */
<1.2.1.3.1.3.1.1.1> /* Colombian government decision on beacon ID/FREQ */
<1.2.1.3.1.1.1.1> /* ID 'R' and FREQ for ROZO on the approach
                     plate correspond with an FMC database entry */
<1.2.1.3.1.1.1.2> /* ID/FREQ combination usually suffice to identify
                     uniquely an NDB within range.
<1.2.1.2.5.1> /* FMC design */
[1.2.1.1] /* CRW turned to "inbound heading" at [1.2.1.3]
             // #DECISION# // #REASONING# */
<1.2.1.2.1> /* CRW unfamiliar with ROZO One Arrival and Rwy 19 Approach */
<1.2.1.2.3> /* CRW used procedural shortcuts */
<1.2.1.2.4.1.1> /* cultural dependencies in ATC/CRW discourse */
<1.2.1.2.4.1.2.1> /* Colombian ATC Use-of-English training/certification */
<1.2.1.2.4.1.3.1> /* no ATC radar coverage in Cali area */
<1.2.1.2.2.1.1> /* lack of time for executing ROZO One arrival procedure */
<1.2.1.2.2.2.1> /* arrival takes place at night */
<1.2.1.2.2.2.2> /* few lighted areas on ground to provide visual reference */
```

Figure 23.6: The Source Nodes