R|V|S
Rechnernetze und verteilte Systeme

Bielefeld University - Faculty of Technology
Networks and Distributed Systems
Research group of Prof. Peter B. Ladkin, Ph.D.

- Home
- People & Info
- Publications
- Research & Projects
- Why-Because Analysis
- Lectures
- Compendium CRICA
- System Safety Society
- Bieleschweig Workshops
- RVS Blog
- Photo Gallery
- Contact & Impressum

Analysing the Cali Accident With a WB-Graph

Thorsten Gerdsmeier, Peter Ladkin and Karsten Loer

Research Report RVS-RR-97-01

Presented at the Human Error and Systems Development Workshop, Glasgow, March 1997.
This work is dedicated to the memory of Paris Kanellakis, who died with his family in the Cali accident.

Abstract: *We analyse the Cali accident to American Airlines Flight 965, a Boeing 757, on 20 Dec 1995. We take the events and states from the accident report, construct a WB-graph (`Why?...Because...-graph) in both textual and graphical form of the 59 events and states, and compare this representation favorably with the Findings section of the original report. We conclude that the WB-graph is a useful structure for representing explanations of accidents.*

Contents

The Accident to N651AA near Cali on 20 Dec 1995

Back to Contents

The accident aircraft, an American Airlines Boeing 757-223, hit mountainous terrain while attempting to perform a GPWS escape manoeuvre, about 10 miles east of where it was supposed to be on the instrument arrival path to Cali Runway 19. 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 FMC (Flight Management Computer) 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 (1). We employ the *WB-graph method* as used in (2) (earlier called the *causal hypergraph* method in (3), (4)).

Our causal analysis compares interestingly with the statements of probable cause and contributing factors in the report.

The Cali Report

Back to Contents

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 critical 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.

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. Each share the logical property that, had they not occurred, the accident would not have happened. We believe it aids understanding to display all such causally-necessary factors and their logical interrelations. The WB-graph, in both its graphical and its textual forms, does that.

The Cali WB-Graph as Formatted Text

Back to Contents

We use an ontology of (partial system) states and events as described in (5). The sequence of events and states used in the graph are those mentioned in the Cali accident report, with one exception. As discussed in (6), 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 (7) also mention fluency training in Aviation English for non-native speakers. As argued in (6), we take the semantics of ATC/pilot English literally, call the affirmation a mistake, and explain this event by the officially-suggested `cultural differences' and `lack of fluency' situations. Readers who disagree should be able easily to make the necessary modification to the graph.

We have found that a path-notation for the causally-relevant states and events is useful. We denote each explanadum as a sequence of digits, e.g., [1.2.1.1] The explanans for [1.2.1.1] is subsequently written as a bulleted list: [-.1], <-.2>, etc, representing the conjunction of all the reasons why event [1.2.1.1] occurred. This is formatted in the form

```
[1.2.1.1] /\ [-.1]
          /\ <-.2>
```

(We use the doubled-symbol `-.' for readability, although one symbol would logically have sufficed.) These explanans nodes for [1.2.1.1] inherit the names [1.2.1.1.1] and <1.2.1.1.2>, respectively. As in (2), we use [...] to denote events, and <...> to denote true state predicates. 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), following the suggestion in (5). Two nodes are classified as both events and states, for reasons to be explained below. These nodes are <[1.2]> and <[1.2.1.3.1]>. We pretty print the nodes according to the lengths of their names. We believe the notation becomes self-evident upon reading. Pilot behavioral failures are classified according to the sixfold classification scheme reproduced here as Appendix 1

```
WHY     BECAUSE                  DESCRIPTION
[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 failure)
               /\ [1.2.2.2.1]
       <1.1.2.1.1> <-.1> CRW unaware of extended speedbrakes
                                               (Awareness failure)
           <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 failure)
     <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]


     <1.2.1> /\ [-.1] CRW turned to "inbound heading" at [1.2.1.3]
                                     (Decision Failure: Reasoning Failure)
             /\ <-.2> CRW without situational awareness: since [1.2.1.3.1.1]
                                                   (Perception Failure)
             /\ [-.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 Failure)
                  /\ <-.4>  PNF didn't correctly verify FMS-entry:
                                  since [1.2.1.3.1.1] (Action Failure)

      [1.2.1.3.1.1] /\ <-.1>      CRW believes 'R' denotes 'ROZO' in FMC
                                              (Awareness Failure)
                    /\ [-.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
                                              (Awareness Failure)
      <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 Failure)
        <1.2.2.2.2> <1.2.1.2> CRW without situational awareness:
                                      since [1.2.1.3.1.1]


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)
```

We constructed the textual form as above, and checking the construction noticed that certain causal factors were missing, namely, [1.1.1.1.1] (*AC dangerously close to mountain*) had no causal forebears. We noticed this discrepancy by singling out the *source nodes* in the WB-graph. These nodes represent causal factors with no causal forebears. Intuitively it should have 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 our method should involve always constructing both.

Automated WB-Graph Construction and Checking

Back to Contents

As noted above, we are aware of the possibilities of error when generating a WB-graph by hand. The first author then wrote 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' (8).

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 take it as a 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 (9)). 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 of us 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 (8) when generating WB-graphs of comparable or larger size.

Event/State Ambiguity

Back to Contents

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. This 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. This affects application of the consistency condition above. We have thus chosen provisionally to classify them as event-states, with the symbol e.g., *<[1.2]>*, and apply the consistency condition formally as for a pure state (that is, not at all).

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]*

To emphasise that it is the `level of granularity' at which the reasons are expressed which engenders this event-state ambiguity, rather than any fundamental problem with the ontology or our method, we note that *<[1.2]>* is very closely related to *[1.1.3]*, *AC very close to mountains @ [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 event-states, 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 event-states.

It is intuitively obvious that a more detailed ontological analysis of flight path/position dynamics will obviate the need for event-states. 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 are not satisfied with event-state introduction because

- it is intuitive: we give no guidance on when to classify factors as event-states, other than to say that it may be expected with position/flight path parameters;
- introducing an event-state relaxes the applicable consistency check on that event-state. 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.

Event-state introduction represents a feature of the WB-method which we wish to develop further, with a view to closer analysis and eventual elimination of the need for event-states in a WB-graph.

The WB-Graph

<u>Back to Contents</u>

Some attempt was made to construct a planar graph. 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) == 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:
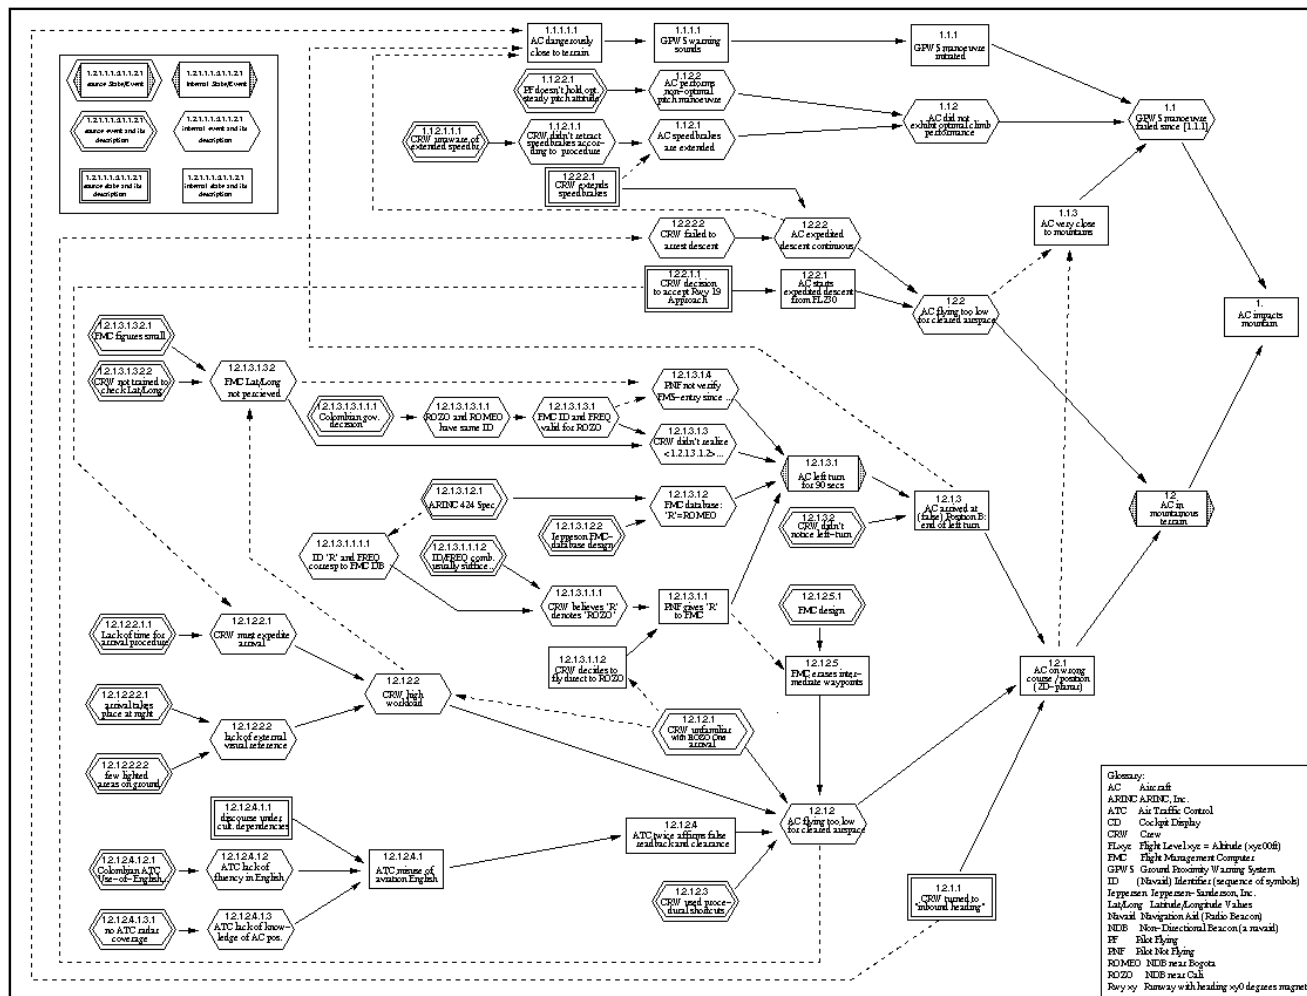
[Link to readable version of the graph](#)

Source Nodes in the WB-Graph

[Back to Contents](#)

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 (`*necessary causal factor* is a transitive binary relation, as noted in (10)) 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.)

```
<1.1.2.2.1>      PF doesn't hold optimal steady pitch attitude
                    in GPWS manoeuvre          (Action failure)
<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 Failure: Reasoning Failure)
<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
```

Critique

[Back to Contents](#)

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 speedbrakes 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 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 (7), 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 for our formal approach to realise that even if an event or state it has reasons in the WB-graph, not all the reasons may be included. The WB-graph method does not (yet) incorporate a method for identifying the important causal events and state predicates. It takes those which have been identified 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 have not included a comparison with procedure in this particular application of the WB-graph method. Such a comparison is necessary, and a technique for performing it is used in the application of the WB-graph method to the `Oops' incident in (2).

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 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 <1.2.1.2.2.2> 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.

Discriminating the `Significant' Events

Back to Contents

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. (11)

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 EHSI (electronic horizontal situation indicator), 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. (11)

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, modifying training requirements to emphasise, for example, paying more careful attention to the EHSI, would not help avoid a repeat; probably this could not cognitively have been accomplished - who knows? 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?

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. We do not solve this problem in this paper.

A Comparison with the Cali Conclusions

Back to Contents

We tabulate and compare the conclusions of the Cali report with the WB-graph analysis. The conclusions of the Cali report may be found in Appendix 2

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 (2), 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 (7). 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 (6).

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.

Conclusions

Back to Contents

We have analysed the causal explanatory relations between the events and states listing in the Cali accident report (1). 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. The WB-method 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.

The WB-graph method 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 loacl 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.

We believe our results demonstrate the usefulness of the WB-method in event analysis.

Acknowledgements

Back to Contents

We are very grateful to Barry Strauch, Chief of Human Factors at the US National Transportation Safety Board, for his detailed and insightful commentary on the first version of this paper, which is particularly visible in the section *Discriminating `Significant' Events.* The paper has been much improved thereby.

We also thank the referees of the *Human Error and Systems Development* Workshop in Glasgow, 19-22 March 1997, where this paper was given, for their helpful comments.

References

Back to Contents

(1): Aeronautica Civil of The Republic of Colombia *Aircraft Accident Report: Controlled Flight Into Terrain, American Airlines Flight 965, Boeing 757-223, N651AA, Near Cali, Colombia, December 20, 1995.* Santafe de Bogota, D.C.-Colombia. Also available at http://www.rvs.uni-bielefeld.de. Back

(2): E. A. Palmer and P. B. Ladkin, *Analysing an `Oops' Incident*, in preparation, to be available at http://www.rvs. uni-bielefeld.de Back

(3): P. B. Ladkin, *The X-31 and A320 Warsaw Crashes: Whodunnit?*, Technical Report 96-08, RVS Group, Faculty of Technology, University of Bielefeld, available at http://www.rvs.uni-bielefeld.de, January 1996. Back

(4): P. B. Ladkin, *Reasons and Causes*, Technical Report 96-09, RVS Group, Faculty of Technology, University of Bielefeld, available at http://www.rvs.uni-bielefeld.de, January 1996. Back

(5): P. B. Ladkin, *Explaining Failure With Tense Logic*, Technical Report 96-13, RVS Group, Faculty of Technology, University of Bielefeld, available at http://www.rvs.uni-bielefeld.de, September 1996. Back

(6): D, Gibbon and P. B. Ladkin, *Comments on Confusing Conversation at Cali*, Technical Report 96-10, RVS Group, Faculty of Technology, University of Bielefeld, available at http://www.rvs.uni-bielefeld.de, February 1996. Back

(7): US National Transportation Safety Board, *Safety Recommendation (including A-96-90 through A-96-106)*, October 16, 1996. Also available at http://www.rvs.uni-bielefeld.de. Back

(8): T. Gerdsmeier, *A Tool for Building and Analysing WB-Graphs*, Technical Report RVS-RR-97-02, RVS Group, Faculty of Technology, University of Bielefeld, available at http://www.rvs.uni-bielefeld.de, February 1997. Back

(9): Ernest Sosa and Michael Tooley, eds., *Causation*, Oxford Readings in Philosophy Series, Oxford University Press, 1993. Back

(10): David Lewis, *Causation*, Journal of Philosophy 70, 1973, 556-567. Also in (10), 193-204. Back

(11): B. Strauch, *private communication*, January 1997. Back

---

Appendices

Back to Contents

Appendix 1: Analysis of Pilot Behavior

To elucidate the pilots' actions, we use an extended *information-processing model*, in which for a given system state, a pilot's interaction with the system is considered to form a sequence:

perception-attention-reasoning-decision-intention-action

This sequence reads as follows.

- *perception*: An annunciation of the system state is presented to the pilot;
- *attention*: the pilot notices the annunciation;
- *reasoning*: figures out what are the possible actions to take;
- *decision*: decides on an action;
- *intention*: forms the intention to carry it through;
- *action*: and finally carries it out.

At least such a fine-grained decomposition of pilot behavior is needed for incident narratives. Failures can occur and have occurred at any stage in this sequence. Examples are:

- During the A330 flight test accident in Toulouse in 1994, an annunciation of the autopilot mode change was not displayed to the pilots, because the angle of attack of the aircraft was higher than 25°. This was cited as a contributing factor in the DGA report.
- In the incident we analyse in this report, the pilot flying (PF) failed to notice that the altitude capture mode was no longer armed.
- In the B757 accident off Puerto Plata, Dominican Republic in 1996, the captain chose to switch on the center autopilot, after concluding that his air data was faulty. The center autopilot obtains its air data from the captain's air data system.
- During the B757 accident off Lima, Peru in 1996, the pilots had lost all effective air data, presumably related to fact that

the left-side static ports were covered with masking tape which had not been removed as the aircraft was returned to service after cleaning. During the incident, the pilot asked for altitude data from Lima Tower, who reported indicating 9,000ft. The PF's AI was apparently reading similarly. He took a calculated risk to begin a descent, and impacted the ocean since his true altitude was a few feet above sea level. His AI read 9,500ft on impact.

- In GLOC (G-induced loss of consciousness) incidents, pilots who regain consciousness are reportedly unable to form the intention to recover an aircraft obviously heading for ground impact.
- The test pilot of the A330 let the departure from normal flight develop, presumably to obtain test data, and initiated recovery too late to avoid ground impact; in the B757 Puerto Plata accident, the crew were unable to take effective action during stick-shaker warnings, allowed the aircraft to stall, and could not recover the stall.

Back to reference

Appendix 2: Cali Report Section 3 (Conclusions)

Section 3. Conclusions (from (1))

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.

Back to reference

Last modification on May 24, 2002

Author(s): Peter B. Ladkin