Air traffic control system failure is too complex to fix in a day

December 22 2014, by Peter Bernard Ladkin

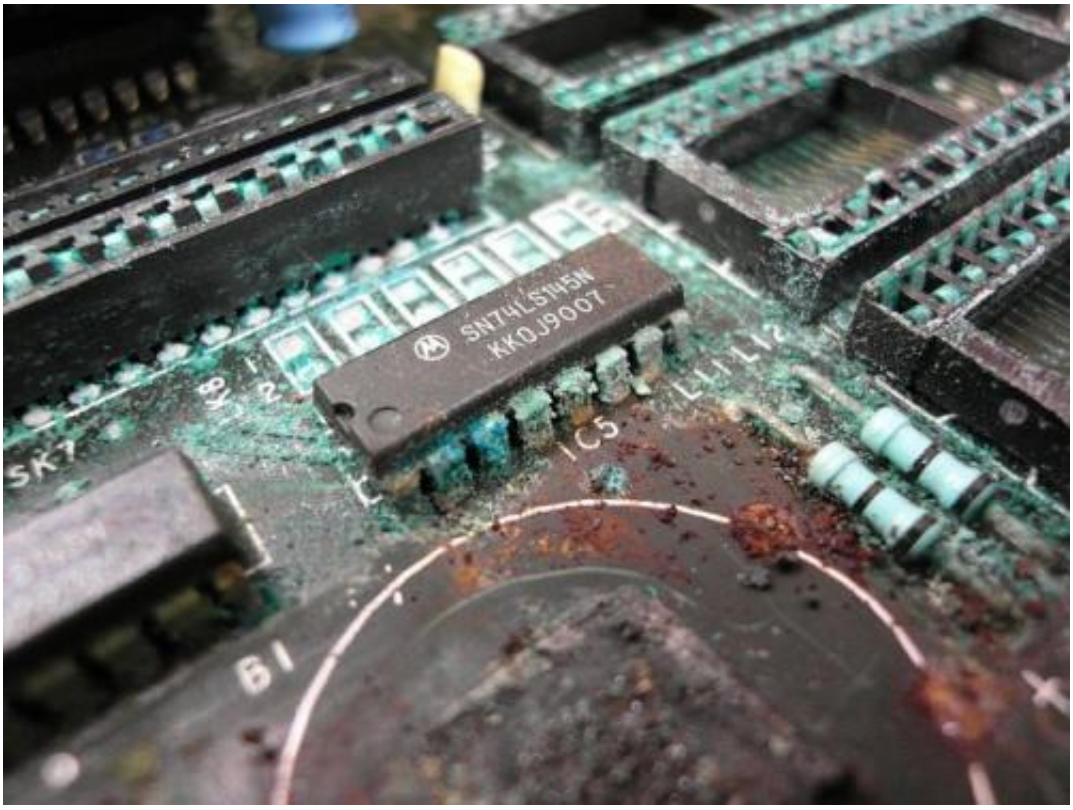

Could air traffic control's ageing, 20-year-old components be to blame? Credit: Binarysequence, CC BY

The recent computer systems failure at the National Air Traffic Services' en-route control centre (known as NERC) at Swanwick in Hampshire led to an airspace slowdown over England and Wales, delaying or grounding hundreds of flights.

The failure lasted 45 minutes – and of around 6,000 flights passing through the affected region that day 120 were cancelled and 500 delayed for an average of 45 minutes. Inconvenient, maybe, but no one was endangered, let alone injured or killed.

Called before the parliamentary Transport Committee, the secretary of state for transport, Patrick McLoughlin, said the failure was "unacceptable". Also appearing at a later session was NATS chief executive, Richard Deakin, who spent some time debating whether "unacceptable" was a term that could be correctly applied in the situation, and if so to whom.

Discussion continued concerning salaries, bonuses, organisational performance measures, the "independence" or not of potential inquiries and how much Deakin worked over the weekend.

I'm a systems person. Calling such a systems failure "unacceptable" is like calling the weather "unacceptable" – nobody wanted it to rain but complaining ain't going to stop it. My questions aren't about salaries or working hours, they are rather: why did the system fall over? Can we expect such things to happen again? Is there is anything anyone can do about it? If so, what?

Inside the problem

The primary failure appeared to be in flight-plan processing, the committee heard, run on a system dating from the mid-1990s. Deakin said the root cause had been identified and a fix put in place to ensure it couldn't happen again.

Now – contrary to worries I expressed to the Transport Sub-committee in 1997-8 during NERC's troubled development – NERC has turned out pretty well, having fallen over only a few times in 13 years of service.

It's inevitable that big, complex, resilient, highly-interconnected programmable-electronic systems such as NERC will fall over eventually.

Some 20-year-old subsystem falls prey to a vulnerability never triggered before, and NATS claims to have discovered the root cause and put in a permanent fix, in just over a day. But hang on a minute. That analyse-and-fix is astonishingly fast for a complex, highly-interconnected system. It suggests to me that the vulnerability was obvious. When aircraft on-board systems suffer such failures, it takes weeks to months to years to analyse – even emergency measures take days to devise. That's because they are subtle; obvious points of failure have already been identified and selected out. Compare:

We were driving down the road, and a wheel fell off. That hasn't happened before. We put it right back on and tightened up all the bolts on all the wheels. It won't happen again.

with:

The air data computer sent a burst of erroneous airspeed spikes to the flight control computer. The flight control computer treated them as true and autonomously commanded pitch excursions [roller-coaster ups and downs], which injured some passengers who were not belted in. We have no idea why those spikes occurred. The flight control computer now filters such bursts out.

This second description is from the inquiry into Qantas Flight 72, conclusions which took years to reach.

No easy solutions

I wonder, is this really a problem solved? It's not easy to devise lasting

solutions to problems that don't potentially bring new problems with them. And if everything is fixed, why is the first item on the terms of reference of the proposed joint CAA/NATS inquiry to review the "root cause"? If, implausibly, there is just one. Almost invariably there are many causes which can be called "root causes", which is part of what makes devising solutions tricky.

When a wheel falls off, in hindsight it's obvious that checking the bolts would have been a good idea. There are engineering methods that prompt us to think of such things in advance which work well for obvious vulnerabilities, but poorly for subtle ones. If Friday's system vulnerability was so quick to analyse and fix, it was likely obvious. So why wasn't it anticipated? What other obvious vulnerabilities are still lying around after two decades? Is anyone looking for them?

I bet NATS has a log of system anomalies which they are working through. How's that going? Are there any gotchas on it which might cause the system to fall over next month? And why didn't anybody at the Transport Committee ask NATS any of this?

Deakin also told the committee that: "We have never seen a repeat occurrence once a fix has been made." That isn't as reassuring as one might think. An investigation by IBM 30 years ago into failures in a big software system estimated that about a third of observed failures would not be expected to arise ever again in the life of the system. That means that, statistically speaking, in about a third of cases doing nothing would be the best solution. By the same token, the same statistics would suggest that, even with perfect interventions, it is only possible to reduce the failure rate by, at most, two thirds.

As Benjamin Franklin might have said, in this world nothing can be said to be certain, except death and taxes and complex-system failure.

Source: The Conversation