Graceful Collapse – Part 2: Common Mode Failure and Resilient Performance in Complex Systems

1. PREVIOUSLY ...

In Part 1 (Owen, 2016), I looked at the contribution of traditional 'predict and withstand' defences in technical systems on their potential for resilient performance. These defences are:

- Redundancy – having more than one system that can perform the same function
- Segregation – splitting a system up
- Diversity – having more than one system that performs the same function, but are different from one another
- Resistance – the inherent capacity of a system to withstand whatever is thrown at it

I looked at how each of these characteristics can contribute to resilient performance. That is, more graceful and less catastrophic failure in the face of unforeseen disruptions and disturbances within the complexity of socio-technical systems.

Now in Part 2, I want pick up on each of these defences again. This time, in terms of their vulnerability to common mode (or common cause) failures. I look at how these failures throw an unwelcome curve ball into the already challenging pursuit of minimising harm in the face of unforeseen disturbances in infernally complex systems. I also look at how common mode failures are like a microcosm of the kind of complex failures that resilience emerged to deal with.

2. COMMON MODE FAILURES AND TRADITIONAL DEFENCES

Common mode failures are as interesting as they are dangerous. They are a particular type of disruption that exploits the physical and functional interdependencies in systems of any type. Think of these as pathways between the elements of a system through which disturbances can travel, cascade and resonate, generating more complex and challenging failures as they go.

Common mode failures happen when a single hazard simultaneously affects multiple systems that are the same, or have very similar attributes, and causes them to fail (Wyman & Johnson, 1997). These types of failures have a way of sidestepping the protections provided by redundancy, segregation, diversity, and even resistance (see Figure 1). They represent a kind of 'beyond design basis' event – failures or combinations of failures that a system wasn't specifically designed to withstand (Winokur, 2012). There is always a potential common mode threat at some level that can cause system failure, and with it, the loss of whatever resilient and other functional capability the affected components provide.

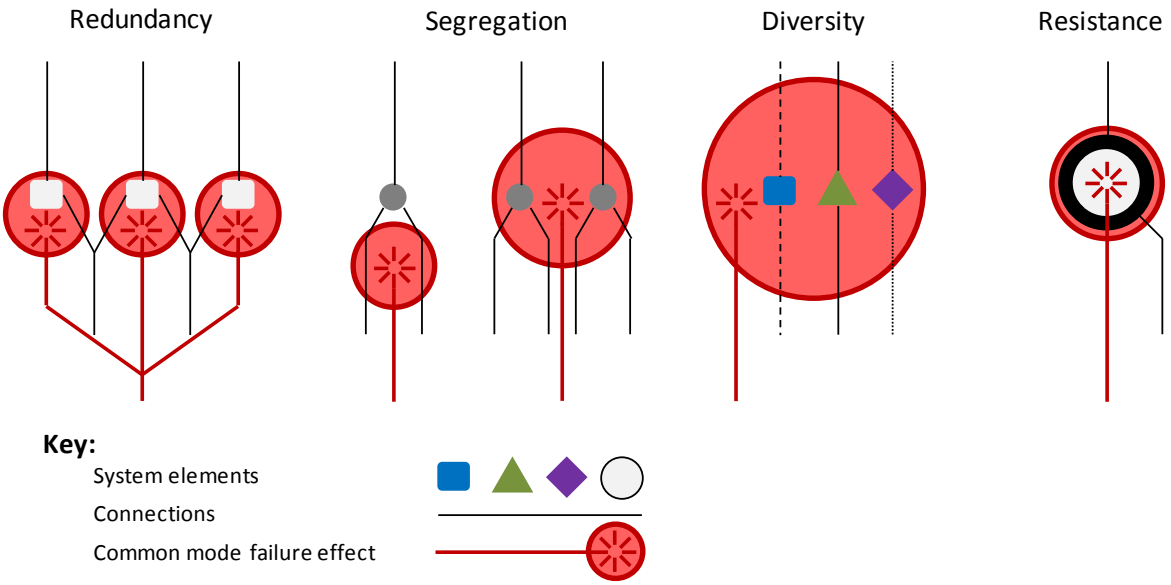

Figure 1. Common mode failure effects on traditional defences in systems

Common mode threats can come from many sources, internally and externally, including from variations in the performance of well-intentioned people in the system. As Reason (2008) noted, '*no system, however sophisticated, is proof against a 'common mode' maintenance error*'.

As a rule, industries don't tend to allow designs with single point failures that can lead to serious consequences. Yet common mode failures have the same impact as critical single point failures by, potentially, affecting all layers of protection simultaneously. This danger can be masked by the appearance of protection offered by implementing traditional defences within our systems. Unfortunately, this sense of security can be misplaced.

Let's start by taking a look at how the protective effect each of the design principles can be undone by common mode failures.

2.1. Redundancy

What is redundancy? Take any four-engine aircraft as an example. Any one of the four identical engines can get us home safely – that's redundancy. And on paper, having more engines makes it less likely that any all of them fail at the same time. But this assumes that failures happen independently. The problem is that they don't, at least not always (Downer, 2009). When we have one cause that affects all engines identically and causes them to fail, that's a common mode failure. It's even worse when they fail simultaneously – that's really bad! Unfortunately, redundancy is perhaps the most vulnerable of all the traditional defences to common mode failures.

What kind of things can cause a common mode failure? A design error and manufacturing fault can affect all engines at the same time. Similarly, a human error in operation or maintenance that is repeated on all four engines can have the same effect. Bitter experience of this very scenario is precisely why we have different maintenance teams, on different shifts, doing engine maintenance on different engines wherever possible (Federal Aviation Administration, 2008).

Our environment is a dangerous place and can be a source of common mode failures too. The now legendary case of the Captain 'Sully' Sullenberger pulling off a spectacular emergency landing on the Hudson River (National Transportation Safety Board, 2010) is a vivid example of both common mode failure and resilient performance saving the day. What happened was the aircraft hit a flock of birds

soon after take-off, disabling both engines. So that was the redundancy provided by two engines done for. At this stage things were looking a bit grim.

But it was the resilient performance of the flight crew that saved the day – a remarkable landing on the Hudson with no fatalities was really down to the flight crew's preparation. Double-engine failure is trained for by flight crews in simulators. In doing so, this serves as a generic protection against any incident causing a total loss of engine thrust, regardless of the specific cause. As I argued in Part 1, generic preparedness to a broad category of failure types can be seen as a resilient capability.

2.2. Segregation

Let's move on to segregation. Physical and functional segregation is effective in protecting systems up to a point. That point is where the scale or strength of a disturbance is sufficient to bridge the divide provided by the segregation between important systems.

Take an example of a shipping company with several offices dotted down the ports of western seaboard of the United States. This physical segregation of offices may protect against localised disruption from things like fire, flood, landslides, pandemic and public disorder. But if a fault line ruptures off the coast causing a tsunami the size of the one on Boxing Day 2004 in the Indian Ocean, that level of segregation may not be sufficient to protect the organisation from being literally wiped out.

Redundancy doesn't help in this case much either, unless other offices capable of providing redundant functions are sufficiently far apart (i.e. segregated) to be outside the tsunami zone. For this to help, the other office must be able to take on the functions of the offices disrupted by the tidal wave well enough to avoid the organisation going under. That requires the people, systems and processes to be in place to enable it to happen, but this is feasible if the investment is made.

2.3. Diversity

On to diversity. Diversity is unfortunately vulnerable to common mode failures in a similar way to segregation. While diversity in design, manufacture and maintenance of systems protects against common mode failures affecting identical systems, it doesn't provide complete immunity. Take Downer's 2009 example of a system with multiple independent power sources or diverse backup systems. These can still be taken out by large environmental or structural disturbances of the kind that affect the structures that the system is attached to. This is a bit like the structural failure of an aircraft or the hull failure of a ship. It doesn't really matter how much redundancy, segregation or diversity you have built into the systems, if the bucket that holds the systems fails, you're snookered.

2.4. Resistance

So this leaves the underlying resistance or strength of the system. Resistance delivers a bit of a margin before things start to go south. However, it can be affected by common mode defects from things like manufacturing, design, maintenance and operational issues too. Any system able to withstand the collapse of the structure it's bolted to with its functionality intact is going to be an extremely rare find.

And let's not forget the people. As well as being a source of common mode threat, the people in our systems are vulnerable to common mode failures themselves through a lack of resistance. Disease, inappropriate training programmes, fatigue, and any number of other factors affecting the

capabilities of more than one person can be seen as a human common mode. This is because they can all reduce the capability of people to do what is needed of them, eroding that precious margin.

3. IMPLICATIONS FOR RESILIENT PERFORMANCE

The examples I've given show that common mode failures represent a kind of 'beyond design basis' event. Traditional approaches are designed to allow a system to take a couple of punches and keep rolling. This provides a limited, but very effective, way to minimise harm and preserve system functions in the face of one or two foreseeable failures. But after that their protective reserves tend to be exhausted. This is typically the result of the cost-benefit equation meaning investment in more layers becomes unappealing, or that adding more layers of bubble wrap just isn't practical.

In the end, it's quite expensive and inelegant to just keep adding more and more barriers and layers of protection. As an elegant response to common mode failures, Hollnagel's Safety II paradigm (2012) is very appealing. It focuses on enhancing our capability for things to go right, and in doing this it has an implicit link to building resilient capabilities.

In terms of getting to grips with the problems that resilient capabilities might help solve, common mode failures are a window into the kind of cascading impact that unforeseen disturbances can have. It's not a pretty sight. But thinking about these effects can help us visualise whether our systems might be able to respond (or not), even if they're not explicitly designed to do so. Because there are typically no specific preparations in place to counter 'beyond design basis' threats, we need to rely on underlying capabilities built into the fabric of a system to get us through – and therein lies the benefit of a resilient capability.

You might still think that the aftermath of common mode failure in an incident depends on luck more than judgement, especially for critical systems. In a way you'd be right. After all, we will be outside of the intended operational parameters. But in another way you might be able to rely on traditional defences within a system that, while not explicitly designed to defend against the disturbance, still have a protective, 'anti-cascading' effect that just hasn't been formally recognised yet.

I think it's fair to say that a system that has this kind of reserve to provide a more graceful collapse in the face of a potentially catastrophic common mode failure is capable of some degree of resilient performance.

4. REFERENCES

Downer, J. (2009). *When Failure is an Option: Redundancy, reliability and regulation in complex technical systems*. Centre for the Analysis of Risk and Regulation, London School of Economics and Political Science.

Federal Aviation Administration, (2008). Extended Operations (ETOPS and Polar Operations), *Advisory Circular No: 120-42B*. U.S Department of Transportation.

Hollnagel, E. (2012). *A Tale of Two Safeties*. (http://www.resilienthealthcare.net/A_tale_of_two_safeties.pdf)

National Transportation Safety Board (2010). Loss of Thrust in Both Engines After Encountering a Flock of Birds and Subsequent Ditching on the Hudson River, US Airways Flight 1549, Airbus A320-214, N106US, Weehawken, New Jersey, January 15, 2009. *Aircraft Accident Report NTSB/AAR-10 /03*. Washington, DC.

Owen, D. (2016). *Graceful Collapse – Part 1: Traditional Defences and Resilient Performance in Complex Systems*. The Schumacher Institute.

(http://www.schumacherinstitute.org.uk/download/pubs/disc/Nov-2016-Graceful-Collapse-Part-1-Traditional-Defences.pdf)

Reason, J. T. (2008). *The Human Contribution: Unsafe Acts, Accidents and Heroic Recoveries*, Aldershot, UK: Ashgate.

Winokur, P.S. (2012). Above and Beyond. DOE Nuclear Safety Workshop, September 19, 2012 (http://www.dnfsb.gov/sites/default/files/Board%20Activities/Board%20Members/Peter%20S.%20Winokur/Speeches/2012/sp_2012919_20356.pdf)

Wyman, R. H. & Johnson, G. L. (1997). Defense Against Common Mode Failures in Protection System Design. International Atomic Energy Agency (IAEA) technical committee meeting on advanced technologies for improving availability and reliability of current and future water cooled nuclear power plants; Argonne, IL (United States); 8-11 Sep 1997.