Graceful Collapse – Part 1: Traditional Defences and Resilience in Complex Systems

1. MINIMISING HARM IN A COMPLEX WORLD IS HARD

Building systems and societies capable of resilient performance is a pretty neat aspiration. It's hard to argue against wanting to minimise harm and maximise success and performance in the face of unforeseen disruption and disturbance. I want to explore whether traditional approaches to defending systems from harm can contribute to resilient performance as well. We'll leave the success side of the coin for another day. The traditional defences can be summarised as:

- Redundancy – having more than one system that can performs 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

This article is about how we figure out what kind of stuff we need to make the best of a bad situation. It's about the choices we can make as designers, commissioners and overseers of socio-technical systems whether they are cities, vehicles, infrastructure, utilities or any other kind of system. And about how we can understand whether the choices we make will have us feeling rightfully smug in the face of unforeseen adversity, or whether there will be an impending palm moving rapidly towards our collective faces at some point in the future.

So once we accept that we live in a very complex world, and we check our rear view mirror for the trail of destruction left by our best intentions and efforts, we can also see that we probably need to do something different to the past. We aren't ever going to foresee everything, and we are going to get knocked by events we never imagined (or at least didn't imagine would happen at the same time that that other thing was going on). We might also be on the receiving end of something so big that there is no way we are ever going to get up, and walk nonchalantly away – a 'beyond design basis' event (Winkonur, 2012).

We must recognise that recovery, or our best approximation of it, is not always on the cards. And that the best we can do is change course, or just try to minimise the fallout. Deep down, we also know that complete system collapse is lurking as an option, and it might take some others down with it. Our world is not always benevolent, but it's in our interests that the demise of our systems is as graceful as possible.

On the flipside, we know that life goes on. We have things to do as individuals, organisations and societies, and we want to do them well and make things better. We'll never achieve that if we spend all our resources trying to wrap ourselves in bubble wrap to protect us every from imaginable threat or error. Pretty soon we run out of bubble wrap, and it's hard to get up and leave the house. It's just not an elegant solution on many levels.

Despite all these collective realisations about the nature of the problem our world presents and its implications, we are where we are. Even though we might recognise that our systems need to do

some pretty radical and funky stuff to deliver resilient performance – things like being reconfigurable, autonomous, and self-organising – we will continue to build systems out of people, processes, hardware and software. These will be an evolution of what we use now and have used in the recent past.

Systems that are reconfigurable, self-organising and autonomous have their own pretty big implications, especially from the perspective of people in the system. But for now, let's agree to accept that these things are at least part of an elegant solution to the problem of complexity at hand.

2. TRADITIONAL DEFENCES FOR TECHNICAL SYSTEMS

Some important things won't change, however. The traditional defensive principles underlying a 'predict and withstand' approach to engineering technical systems will surely remain. These make our systems capable of more graceful and less catastrophic failure into of redundancy, segregation, diversity and resistance providing greater margin.

And why not? These guys have an enviable track record of keeping fiery infernos of death and calamity at bay for both technological and human elements of systems. As these are already well established principles employed in the design of many safety-critical systems, it makes sense to try to understand what contribution they can and can't make to resilient performance in terms of minimising harm and maximising performance.

In Downer's excellent 2009 paper, he discusses the issues especially around redundancy, diversity, independence and complexity in some detail, so I won't repeat that here. Instead, what I will look at is the potential contribution of these design characteristics to resilient performance, and the people within the system. I have added 'resistance' to this list because it's important to capture that the underlying strength of systems to survive the forces they are exposed to contributes to survivability in a similar way. This is sometime known as 'margin', and in resilience terms can be thought of as an ability to resist disorder (Fiskel, 2003).

3. CAN TRADITIONAL APPROACHES DELIVER RESILIENT PERFORMANCE?

Before we get going, it's worth flagging that the defensive principles of redundancy, segregation and diversity bring with them an uncomfortable paradox when trying to develop systems capable of resilient performance. They have a tendency to add complexity to our systems (Downer, 2009). So the characteristics that we might look to to develop the resilient capabilities of systems might be also be adding to the problem we set out to solve, i.e. impenetrable and unpredictable complexity. *C'est la vie*.

There are a couple of important aspects that I will leave for another day. In Part 2, I will look at common mode failures, as they have a nasty habit of unravelling our best efforts to protect our systems using traditional defences. And in Part 3, I will look deeper to understand whether the capability to change the level of each type of defence on the fly has potential benefits for resilient performance (e.g. changing the amount of segregation).

Let's look at each type of defence in turn, and see how it stacks up in building a resilient capability. This is summarised in Table 1.

Table 1. Summary of traditional defences, their resilient capabilities and vulnerabilities

Defence	Potential Resilient Capability	Additional Complexity*	Vulnerability
Redundancy	• Allows reconfiguration using redundant systems • Can (temporarily) increase capacity by bringing redundant systems on line	+ +	• Capacity may be degraded depending on performance of redundant systems where one or more has failed • Reduced overall reliability by increasing wear across the system
Segregation	• Defence against unforeseen threats of the same generic energy type (e.g. Kinetic, electrical, thermal, etc.)	+ +	• Connections add failure cascade pathways • May only provide defence against specific threat types
Diversity	• Allows reconfiguration using diverse systems less vulnerable to common mode failures	+ + +	See 'Redundancy', plus: • Increase in system complexity increasing possibility of unanticipated and emergent • Complexity hampers usability and sustaining situation awareness
Resistance	• Retain margin of performance over a wider range of operational conditions resisting effects of volatility	+	• May only provide margin against specific types of threats • Can limit reconfigurability

*Key: + None / low + + Moderate + + + High

3.1. Redundancy

In technical systems, redundancy is having another identical system. This means that when one fails there is another to do its work for it, and maintain the function of a system within acceptable limits. Having four identical engines on an aircraft provides redundancy. You only need one to fly so we can afford to lose three and still limp home. Another example is two water mains serving a city. If one fails, becomes disrupted or contaminated, there is another system to pick up the demand. Whether the overall system continues to function at the same level or not depends on the capacity of the remaining redundant components.

In terms of resilient performance, redundancy potentially allows for some degree of system reconfiguration. We can switch between systems, turn them off or on. In a city we can close or open roads to redirect traffic flow around an accident. We can also increase capacity by bringing redundant systems on line and operating them simultaneously. Motorways can be managed this way by opening the hard shoulder to cope with additional traffic during peak times. However, operating redundant systems simultaneously could impact the overall system reliability by increasing wear across all parts of the systems.

It also doesn't really matter what specifically has taken out the main system – it can be something foreseen or unforeseen, either way redundancy gives us options.

While there is another path, the overall functional capacity (i.e. margin) could be reduced, so once we've engaged redundant systems we are running a bit closer to the wire. There will be a finite number of failures before we run out of redundant systems. But while harm may be done in terms of throughput, it won't be as debilitating as relying on a single solitary system that could fail, which could be catastrophic.

3.2. Segregation

Segregation is a way to prevent failures cascading though a system. Cascading failures are really bad news because they cut a swath through interconnected parts of systems that we'd really prefer stayed on speaking terms. Segregation allows some parts to fail without taking down other parts of the system – it preserves their independence.

We build in segregation in two ways: we physically separate the location of systems that do the same thing (e.g. you could have half your IT team in India and half in America), and we can have separate feeds going into and out of systems (e.g. two electrical power lines to a pump from two different generators). Using both of these approaches together is a win-win: they stop the failure of a system from affecting other parts, and the failure of other parts of the system affecting it. Neat.

The idea of connectedness is important in resilient performance. To act as a system, at some level a system's parts will be connected, acting in concert to achieve a goal. But connectivity is a double edged sword – connections act as pathway for disturbances to transmit down (e.g. like catching a cold from colleagues in a meeting), but they also act as ways to reconfigure and transmit resources through alternative (redundant) routes.

The value of segregation for unforeseen events is a bit hit and miss. We typically segregate systems in response to specific threats. Aircraft engines are physically separated to prevent the failure of one engine spitting broken fan blades into the one next to it and causing it to fail too. On the surface, it's difficult to imagine how to design in segregation without first foreseeing the type of threat.

But if we dig deep enough, we can find a useful way to understand threats in different broad classes. Not wanting to get too high-school physics, but we can group threats together by the type of energy that does the damage. So we can understand what our segregation achieves by what it kind of energy it protects the system from, e.g. kinetic, thermal, radiological, electrical, chemical, etc. People are affected by these energies too, but we might also want to add pathogens to that list. Which, curiously enough, also affects software through viruses.

So the value of segregation to resilient performance in the face of unforeseen threats is there. We may have foreseen the impact of an earthquake (a kinetic disruption) on an accounting department located in New Zealand, and in doing so we may have also catered for a truck crashing through the wall of the office, but without having explicitly identified it.

3.3. Diversity

In terms of resilient performance, diversity looks like the mother lode. It's like redundancy turned up to 11. Not only do we get different systems able to take on the work of others when they fail, but

each of them is different (i.e. diverse), so they are less likely to be affected by one type of disturbance.

Another example from aircraft: typically three braking systems are used – normal, backup and emergency. Each is designed and manufactured by a different company to make sure that the way it works and the manufacturing process are different. Functional diversity goes even further using different methods to achieve the same goal. For example, one sensor could use pressure, and another uses temperature (Downer, 2009).

In terms of people, diversity is of benefit too. Different backgrounds, skills and approaches are incredibly valuable to avoid groupthink (Tetlock, 2015) and to generate innovative solutions to novel problems.

So we're good then. For resilient performance we should just build systems with lots of diversity, right? Well, if there is any virtue that adds significant complexity to a system, it's diversity. So, yes, we should definitely embrace diversity as a means of achieving resilient performance, but we should be watchful for over-complicating our systems through adding diversity (or redundancy and segregation). This adds to their unpredictability and to the opacity of system behaviour and failure, and it also amps up the challenges of use, operation and maintenance every day and during emergencies.

But, on a positive note, there appears to be no real downside to involving groups of diverse people, providing they are carefully assembled to maximise difference.

3.4. Resistance

This last characteristic is a bit more fundamental. It's what we used to do before we invented redundancy, segregation and diversity as more sophisticated ways to dodge disaster. We make things stronger, harder – more able to resist the pressures, temperatures, vibrations and stresses of life.

Like all the other characteristics, hardening a system can work for and against us, especially when we are trying to manage events that are beyond the design limits of our systems. Let's take the positive side first.

Hardening a system can mean it can take more punishment, or operate to greater extremes before we see a reduction in performance, or we experience some kind of harm. With the recognition that complexity can bring greater volatility (e.g. climate change bringing greater extremes of temperature), this is a simple and effective way to give us some leeway, and doesn't add to the complexity of our system. For instance, building higher flood defences gives us more margin in terms of the height of flood water. Depending on the system, we may also be able to increase this limit even further through reconfiguring the system by adding sandbags to the top of the flood bank to extend their operational envelope, albeit temporarily.

However, depending on the nature of the system, the technological approach behind hardening may only be effective against some types of threats (e.g. heat but not pressure). In addition, full credit could only be given to a system that also resists a rapid rate of change in environmental conditions without becoming brittle.

The problem can be that some systems are very difficult to give more resistance to because of how they are made. It's very difficult to give a nuclear reactor a bit more thermal resistance, during an

impending meltdown, with a roll of duct tape and the office ice machine. It is kind of interesting to think about whether it is useful to build in 'bodge-ability' to high tech systems, to be modified on the fly, and what that might look like in reality.

Similarly, having too much hardness in our systems can be problematic. Let's take the reasonably common case of extreme flooding events where the flood banks have been overtopped and water has inundated areas we really wanted to keep dry. Sometimes we want to break through flood defences to let water back out. For earthen defences this can be done quite easily with a digger (or a few shovels), but if the defences are made of reinforced concrete this is no longer an option. A concrete system is tough, but fragile at the same time, and hence not so capable of resilient performance in that way (but yet may have more margin – yep, it's complicated!).

4. WHERE THIS LEAVES US

The good news is that despite their limitations, traditional approaches to defending systems against foreseeable hazards appear to, in principle, offer some resilient capability by providing a degree of reconfigurability. That's great.

Overall, I believe that it's really important to recognise the aspects of existing systems that already support some level of resilient capability so we make sure they don't get changed or dismantled without understanding the consequences.

So more redundancy, segregation, diversity and resistance are good things to a point – that point being when the system becomes so complex that its complexity is a hazard in itself. But we can't assume that a system is automatically resilient just because it has these defences. If a system includes redundant elements, but is not able to be reconfigured by the operators or automation, then its contribution to resilient performance will be lost.

In reality, systems don't fit neatly into one or the other category of defence, but that's OK. We should be able to give credit for all resilient capabilities, but may only rely on one in a given scenario. For example, I might have access to two cars so that if one doesn't start I can get to work as normal. While the cars are diverse because they are made by different manufacturers at different times, here I'm leaning on the redundancy more than the diversity to keep me out of trouble with the boss.

The bad news is none of these approaches are impervious or insurmountable to the complex array of threats and interactions. That's perhaps to be expected. But it's useful to understand what kind of vulnerabilities our systems have to different kinds of threat, especially if we can do this for generic types of threat.

This generic approach sits in a pragmatic grey area between foreseen and unforeseen threats. This is useful when thinking practically about what system characteristics can help achieve some degree of graceful failure through the application of resilient thinking. So while we are unlikely to foresee all specific threats, we can at least understand what any kind of threat that delivers a certain type and magnitude of disruptive energy, and what this would do to the functionality of our systems.

I will admit that this falls short of being a solution to completely unforeseen events, but we would have taken a step in the right direction and it's certainly more elegant and efficient than a brute force approach that tries to foresee everything at an immeasurably enormous price. After all, there isn't an industry out there that delivers safety at any cost. There is always a limit, and seeking resilient performance is a way of being more effective within it.

5. 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.

Fiskel, J. (2003). Designing Resilient, Sustainable Systems. *Environmental, Science and Technology*, *37*, 5330-5339.

Tetlock, P. & Gardner, D. (2015). *Superforecasting: The Art and Science of Prediction*. Random House Books, UK.

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).