Resilience is an Emergent System Property: A Partial Argument

Peter Bernard Ladkin, Bernd Sieker

Causalis Limited & Causalis IngenieurGmbH
Bielefeld, Germany

Abstract *Systems are collections of objects exhibiting joint behaviour. Sometimes this behaviour is anticipated, sometimes not. We have studied a number of types of complex systems and their failures, including electricity supply grids, motorways, the financial system, and air traffic control. We argue that the resilience properties of such systems are largely emergent. We illustrate the thesis through analysis of three electricity blackout events. We consider one event in detail and two others summarily.*

1 Systems

In order to talk about properties of systems, including emergent properties, we first need some definitions. We have been using the following definitions for over a decade (Ladkin 2001, Chapter 3). We will not use here the formal properties of these definitions, but we judge it is well to state the vocabulary and its meaning to us. We illustrate the definitions below.

- A system is a collection of agents with joint behaviour
- An agent is an object with behaviour
- Agents have properties (attributes)
- Multiple agents have relations
- Behaviour is considered as: change in properties and relations over time

Systems have boundaries: some agents and other objects belong in a system; others are outside. Natural system boundaries are often drawn to satisfy the following criterion (note that this is just one criterion; there are others): relations/joint behaviour of objects that “crosses the boundary”, that is, some objects in the joint behaviour are in the system and others out, are relatively sparse, whereas the relations/joint behaviour of objects, all of whom are in the system, are relatively abundant.

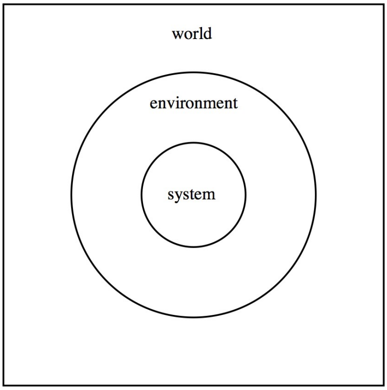

Figure 1: A Venn Diagram of a System, Its Environment, and the World Outside (Ladkin 2001)

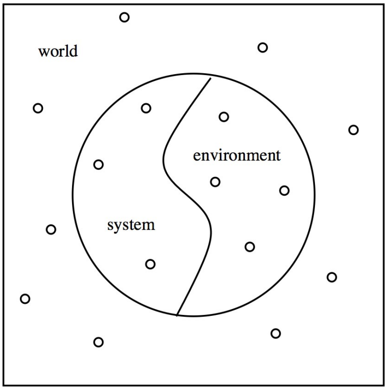

Figure 2: A System May Interact with Objects Not in the Understood Environment (Ladkin 2001)

When a system is conceived, and its boundary is drawn, the system designers try to understand the environment as all objects outside the system with which the system interacts; which the system influences or vice versa. The ideal is shown in a Venn Diagram in Figure 1. However, mistakes may well be made. There may be parameters simply missed; interactions not foreseen or understood. The reality is more often as in the Venn Diagram in Figure 2.

When a system is conceived and designed, various parameters, properties and relations of system objects and environment, are laid down and the behaviours specified. This may be formal or, more usually, informal. A state of a simple system, formally described, is shown in Figure 3.

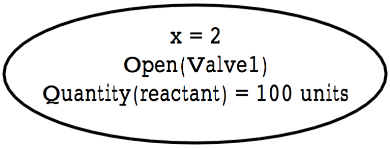

Figure 3: A State of a System, Formally Described (Ladkin 2001)

A behaviour, a change in system state, is shown in Figure 4: Valve1 is opened; nothing else is changed - parameter x and the quantity of reactant retain their values.

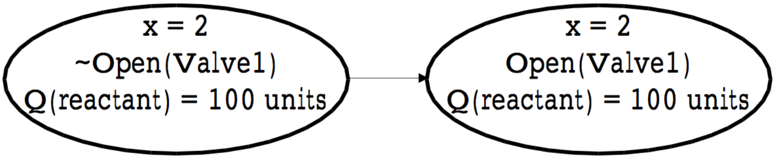

Figure 4: A Behaviour (Ladkin 2001)

We shall use the following conception of *emergent property*. Emergent properties of a system are properties which are not base; that is, they are not defined in terms of the properties and relations of the objects constituting the system, as designed or conceived.

Emergent behaviour is joint behavior of objects which is/are (behaviour/objects) not initially considered. One kind of system of systems which is currently attracting a lot of attention consists of swarms of small, simple aircraft. Swarm behaviour cannot be completely described using the individual properties and relations of a single agent or its immediate interactors. A new vocabulary is often needed; the swarm behaviour is thereby emergent. An example is murmurations of starlings, shown in Figure 5. Describing the density and movementof the mass of starlings, it may well be that the vocabulary of fluid mechanics would be useful, which is far away from any vocabulary useful for describing individual starlings or their interaction with neighbors.

One often overlooked way in which properties can turn out to be emergent is when there is a failure scenario of a type which had been unanticipated. This very often happens when the hazard analysis of the system is incomplete. Here is an example. Hazard analyses are often conducted with the help of techniques such as FMEA1. The FMEA conducted on the Boeing 787 Lithium-ion-type main and auxiliary batteries considered the phenomenon of thermal runaway. The Boeing analysis identified overcharging as the only event which would result in (smoke and) fire. An FMEA by the battery manufacturer GS Yuasa suggested that an internal short-circuit would only result in smoke production (NTSB 2014, pp50-52).

We understand the analysis was not revised when a battery under test underwent thermal runaway and burnt down a building in November 2006 (*op.cit.* p43). Batteries ignited twice in 2013 on Boeing 787 aircraft in line service and the probable cause of one of the incidents was "*internal short circuit within a cell of the auxiliary power unit (APU) lithium-ion battery, which led to thermal runaway that cascaded to adjacent cells, resulting in the release of smoke and fire"* (*op.cit.* p79).

Figure 5: A Murmuration of Starlings
(© Walter Baxter, reused under a Creative Commons Licence)

We consider the resilience of systems of electricity supply through the grid. In other work, we consider collision avoidance in air traffic operations, collision avoidance in rail operation, auffahr accidents on motorways, and asset protection and enhancement in banking and company activity.

We call systems of the sort we consider *teleological systems*. They are systems built (by people or other animals) for a purpose. Teleological systems are distinguished from such naturally-occurring systems such as predators and their prey, or other ecosystems, which have system characteristics but no overt purpose intended by any conscious entity. Most engineered systems are teleological. Some systems are not. For example John Conway's Game of Life (Conway 1970) is a mathematical system whose original purpose was, if anything, for its creator and others to have fun. One emergent property of the Game of Life is its usefulness in illustrating the talk accompanying this paper.

We need a definition of resilience. The convenor of the EU ReSIST project[1], Jean-Claude Laprie, defined it as "*The persistence of service delivery that can be justifiably be trusted, when facing changes*", cited in (Meyer 2009). Meyer also considers the definition "*the ability of a system to deliver service under conditions that lie beyond its normal domain of operation,*" as well as others, such as that of

[1] Full Disclosure: The first author was a formal reviewer of the ReSIST project.

Woods: "*how well can a system handle disruptions and variations that fall outside of the base mechanisms/model for being adaptive as defined in that system.*" (*op. cit.*)

2 Electricity Blackouts

To understand how blackouts may happen, it is necessary to understand some qualitative physics of grid supply. Except for a few direct-current (DC) lines, almost the entire European grid is a synchronised, in-phase alternating-current (AC) grid. The North American grid system is a set of three grids, with some inter-grid connection through high-voltage DC lines. The main reasons for using AC are historical, as alternating voltage conversions are technically trivial using transformers.

The downside of AC is that the entire connected network must be at exactly the same frequency and in phase to avoid large energy losses. The frequency may change as a result of the mechanics of energy – see below. Such change must be actively managed.

In most power sources (power stations), mechanical energy is converted into electrical energy. In some, nuclear energy is converted into heat and via steam and mechanical energy into electrical energy. Large centralised power stations include nuclear, coal-fired and gas-fired stations, and hydroelectric stations. Decentralised power generation includes wind turbines (mechanical into electricity) and photovoltaic installations (light into electricity).

AC electricity supply divides into active power and reactive power. Instantaneous power is, as in other areas of physics, the product of voltage and current at an instant. These quantities vary sinusoidally with time in AC supply. "*Sinusoid*" means the following, illustrated in Figures 6 and 7. Suppose there is a circle with centre at (0,0) on a two-dimensional surface, and a radius of that circle which is rotating at constant angular speed. Then the sinusoid quantity is the y value of the (x,y) values traced out by the tip of the radius. The angle which the radius makes with the x axis is called the "*phase*" of the sinusoid. In Figure 6 the radius is given by the blue line labelled U_C.

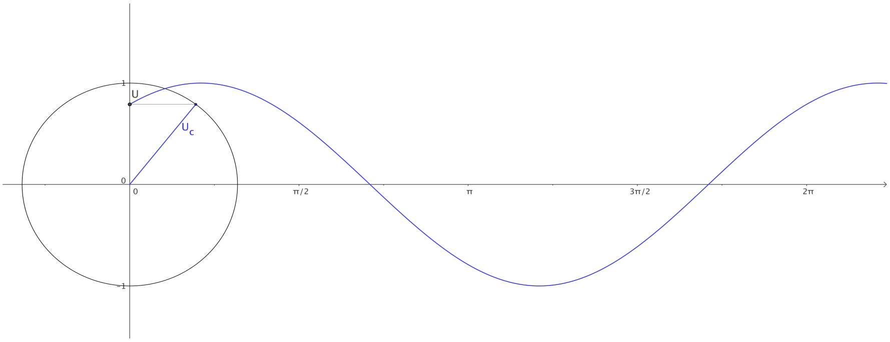

Figure 6: A Sinusoid Curve generated by Voltage over time

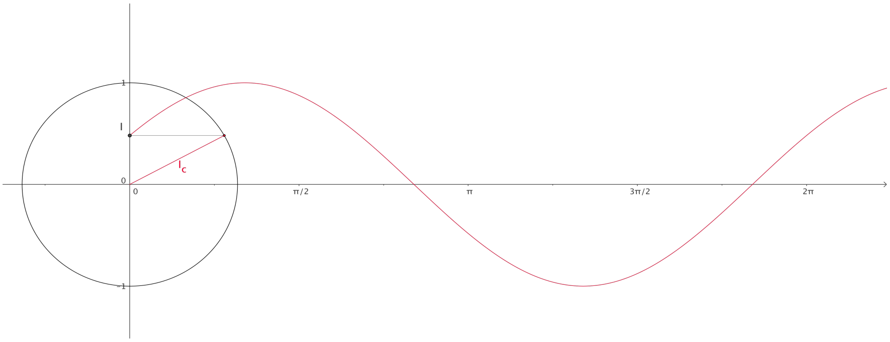

Figure 7: A Sinusoid Curve generated by Current over time

Voltage and current are two parameters of electricity flow which result in sinusoid fashion from AC generators. Their respective values given by the sinusoid curves may not cohere.

Power is the product of voltage and current. When voltage and current are "*in phase*", then their values are always either both zero or both positive or both negative, so their product is zero or positive and so the power delivered over one cycle (= rotation of the radius) is the integral of that and is positive. That power may be used to do work in a recipient device and is called "*active power*". This is illustrated in Figure 8. Instantaneous power is shown by the brown curve, and its mean over a cycle by the area under this curve, here normed to 1.

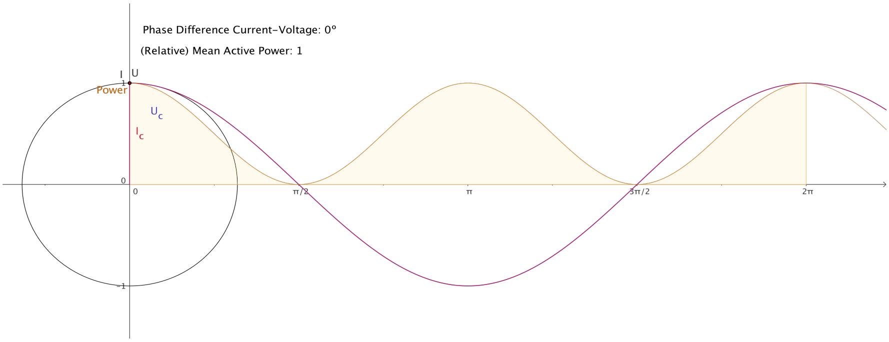

Figure 8: “In phase” voltage and current – positive average power

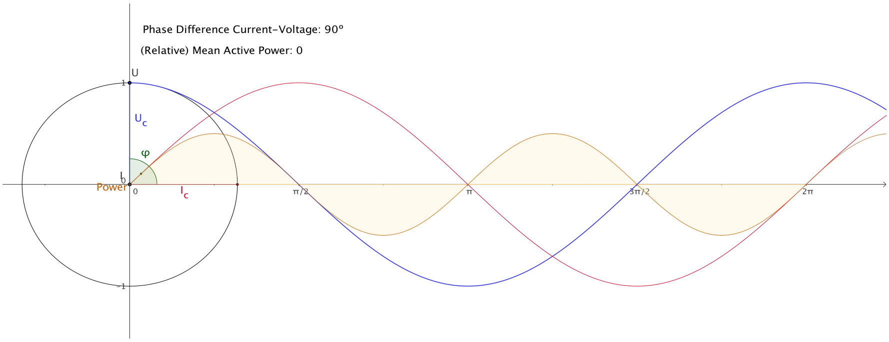

Figure 9: 90° “out of phase” voltage and current

If the voltage and current are 90° “*out of phase*”, then the integral is zero over one cycle. This means that the average power delivered over a cycle is zero. Such an out-of-phase current is called “*reactive power*”. This situation is shown in Figure 9. The general situation is given in Figure 10.

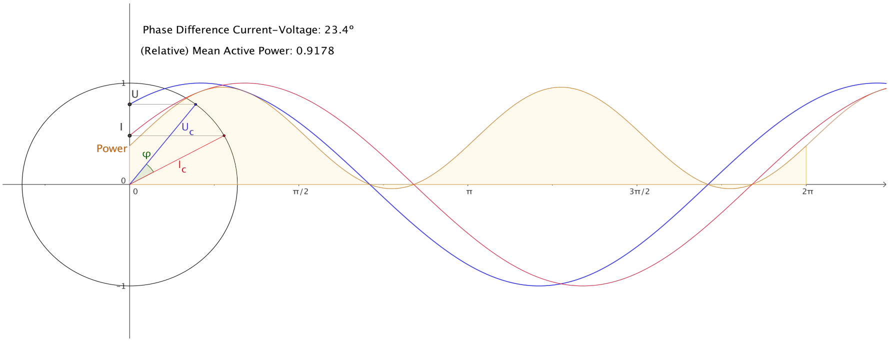

Figure 10: General case of “out of phase” voltage and current

Any phase is mathematically given by a vector sum of orthogonal components. It follows that any AC power, as in Figure 10, can be considered as a sum of active and reactive power. Relative phase shifts, and thereby the reactive power component, is caused by capacitive and inductive elements in the grid. Reactive power has a large influence on voltage levels and creates additional losses on the transmission lines. It must therefore be carefully controlled.

Energy is also stored in the electrical grid in other ways. For example, surplus power in the grid can be absorbed by generators if their drives are disconnected – they turn into motors and electrical energy can be turned into momentum of the rotating armature. When this happens, the armature will speed up, so that when the item is reconnected as a generator its frequency is a little higher[2]. Conversely, the spinning armature can temporarily supply additional energy into the grid, by transferring some of its momentum into electric power, thereby slowing down. Such frequency fluctuations (of the nature of mHz to, say, cHz) are appropriate over a short period of time, provided they are counteracted before the phases difference to adjacent parts of the net becomes too large.

The reliability of a grid-based electricity supply means providing electrical energy at the required voltage and in the demanded amount to consumers. "*Maintaining reliability is a complex enterprise that requires trained and skilled operators, sophisticated computers and communications, and careful planning and design.*" (US-Canada Task Force 2004).

2.1 The 2003 North-Eastern North American Blackout

We consider first the August 2003 blackout of large portions of the Eastern Interconnection in the USA and Canada. The North American grid is divided into three, each component called an "Interconnection". An interconnection is a more or less open network in which the flow of electricity is physically determined by supply and demand, operating according to basic laws of physics. Flow can be controlled only be regulating supply and demand. Within an interconnection, there are usually many pathways available to satisfy a demand, and this yields a certain resilience. Connections between the interconnections are often established by DC lines. The North American interconnections are shown in Figure 11.

[2] Large power stations usually employ synchronous generators where the AC frequency is tied to the rotational speed of the armature.

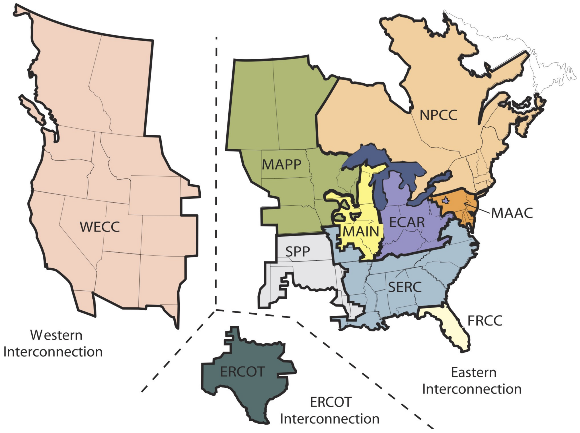

Figure 11: The North American grid interconnections (NAERC 2014, under fair use)

August 14, 2003 was not a particularly hot day in the U.S. Midwest, with temperatures in the mid-80's Fahrenheit (28-30°C). Electric air conditioning systems were operating in many homes, but the electricity distribution system (the "grid") had dealt with 100°F (38°C) temperatures a year before without problems.

Within each interconnection, power supply and demand must be matched, else frequency fluctuations occur which can damage equipment. Reactive power must be balanced to maintain acceptable voltages; voltage fluctuations can cause a collapse in supply when low, and can damage equipment and result in arcing when high. Electricity flow over transmission lines heats up the lines and must be controlled to ensure that thermal limits are maintained; hot lines expand and sag, and clearances from other objects must be maintained. Furthermore, flows must be managed to absorb "contingency events", such as a generator going off-line or a transmission line "tripping" (shutting down). External insults such as contact with trees or physical damage to lines are usually handled through tripping.

The North American Electric Reliability Corporation is a voluntary organisation whose mission is to assure the reliability of electrical power supply in North America. Its members are ten regional reliability organisations. The August 14 blackout affected three of the ten regions. There are 140 "control areas" in the US. A control area has one entity, either an "independent system operator" (ISO) or "regional transmission organisation" (RTO), responsible for balancing generation and loads in real time to maintain stability (one of two primary functions determined by legislation). They also control generation directly, to support interchange schedules with other control areas, and operate collectively to maintain stability of their interconnection. Control area dispatch centers monitor and control electricity generation and flow and are staffed continuously.

The initiating events of the blackout involved two control areas, FirstEnergy (FE) and American Electric Power (AEP), and their reliability coordinators, Midwest Independent System Operator (MISO) and PJM Interconnection (PJM). FE operates a control area in northern Ohio. FE consists of seven electric utility operating companies, four of which operate in the NAERC ECAR region, with MISO as their reliability coordinator. AEP operates a control area in Ohio just south of FE. AEP is both a transmission operator and a control area operator. PJM is AEP's reliability coordinator.

The area of blackout is shown in Figure 12. The two-nation Task Force Report (U.S.-Canada Task Force 2004) identifies four classes of causes of the blackout:

1. FirstEnergy and ECAR failed in general to assess and understand the inadequacies of FE's system, particularly concerning voltage instability and the vulnerability of the Cleveland-Akron area. FE did not operate its system with appropriate voltage criteria.
2. There was continuing inadequate situational awareness at FirstEnergy. FE did not recognize or understand the deteriorating condition of its system.
3. FE failed to manage adequately tree growth in its transmission rights-of-way. Tree contact was the cause of the outage of three FE 345-kV transmission lines and one 138-kV line during the incident.
4. The ISO/RTOs failed to provide effective real-time diagnostic support.

It is interesting that this enumeration of causes, as well as a similar but lengthier list in the NAERC report, concerns exclusively human or organisational failures. The physical causes of the event, what a physicist or scientist would say were the causes, are omitted, as well as some system characteristics which appear to us to be crucial in explaining system behaviour.

The detailed history of events (NAERC 2004), shown in general area form in Figure 13, enables further observations about how the grid system functions.

- As we have noted, an interconnection is a network flow operating under purely physical laws. Stability of flow is maintained through partly automatic and partly human intervention. Specific loading (energy consumption by consumers) varies according to circumstances generally not under network control. Network controllers can moderate active power flow; also reactive power flow in order to equilibrate its generation inside the network. Methods for moderation include supplementing generating capacity (increasing power output from generation plants, or bring off-line generation equipment on-line). In principle, transmission lines may also be taken out of service, but this did not happen in this event. Transmission lines took themselves out of service ("tripped", then "locked out") for a variety of reasons.
- A, maybe *the*, key event in this blackout process was the tripping of the Sammis-Star 345kV transmission line at 16:05. This is shown in Map 3 of Figure 13. This "*completely severed*" (*op. cit.* p55) the 345kV transmission path from South-eastern Ohio into Northern Ohio (Cleveland-

Akron, on the southern shore of Lake Erie, and Toledo on the western shore), which had significant net import of power at the time due to the demand. Three pathways were still available, namely from northwestern Pennsylvania to northern Ohio around the south shore of Lake Erie, from southern Ohio, from eastern Michigan and Ontario. However, "*no events, actions, or failures to take action after the Sammis-Star trip can be deemed to have caused the blackout.*" (*op. cit.*, p 56). In other words, in the demand-supply circumstances prevailing at the time, the Sammis-Star line trip was a single point of system failure. At the point of tripping, the reactive power carried on the line was ten times as high as earlier in the day.

- Previous to the Sammis-Star line trip, the NERC report suggests that operator load-shedding may have been appropriate to maintain stability of the system but that afterwards "only automatic protection systems" would have mitigated the consequences (*op. cit.*, p57).
- The cascade developed into a blackout for "three principal reasons" (*op. cit.*, p58):
  - Loss of the Sammis-Star line triggered many subsequent line trips;
  - Many lines operated with so-called "zone 3" impedance relays, as did Sammis-Star, which respond to overloads rather than line faults;
  - Relay protection settings for line, generators, and under-frequency load shedding may not be sufficient to reduce the likelihood and consequences (= risk) of a cascade, "*nor were they intended to do so*" (*op. cit.*, p58).

The blackout happened within about seven minutes after the Sammis-Star line trip. There were large power surges, for example a 3,700 MW flow from Michigan to Canada turned into a 2,100 MW flow in the other direction within one second, a 5,800 MW flow reversal. The events caused a large electrical island separated from the rest of the Eastern Interconnection. The region had been importing power and did not have enough generational capacity within to satisfy the demand. However, pockets within this island did stabilise, and recover. Phase and synchronisation mismatches often hinder facility resetting after trips. Some pockets took a long time to recover.

The reason for the Sammis-Star 345kV line trip was that a protective relay sensed "*low apparent impedance*", that is, low voltage and high current (*op. cit.*, p57). There was in fact no fault. The protective relay cannot physically distinguish a fault from high load, and the line was operating at 130% of nominal capacity and the voltage was lowering. As mentioned above, before this event operator load-shedding could have reduced load, thereby avoiding the trip, but after the event "*only automatic protection systems would have mitigated the cascade*" (*op. cit.*, p57) and there were none, or insufficient, in place.

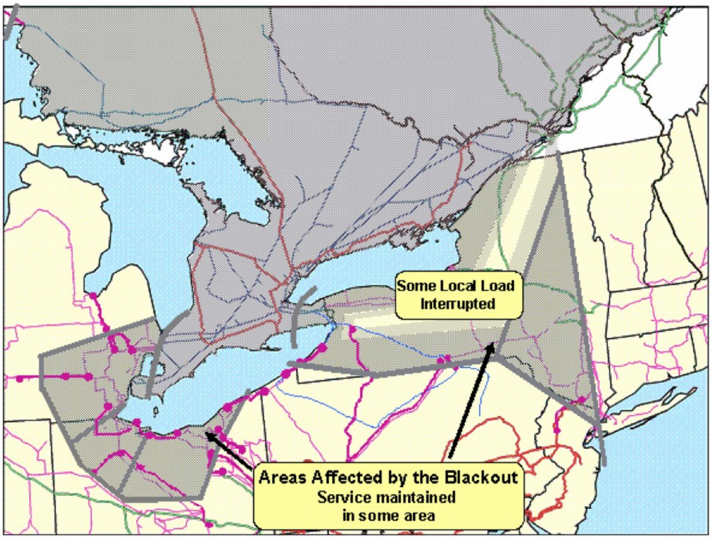
Some Local Load Interrupted
Areas Affected by the Blackout
Service maintained in some area

Figure 12: Areas affected by the 14 August 2003 blackout
(NAERC 2004, reproduced under fair use)

The above observations concerning the Sammis-Star trip can be summarised as:

- The Sammis-Star trip led inexorably, under the supply-demand conditions prevailing at the time, to blackout;
- The Sammis-Star trip was a normal, designed reaction to the conditions on the line, namely a condition of “low apparent impedance”.

Two conclusions follow directly from these observations:

- Given the conditions prevailing at the time, the Sammis-Star line was a single point of failure of the grid;
- The blackout as a consequence of the Sammis-Star trip was a “normal accident” in the sense of Perrow (Perrow 1984). It was an undesired event which followed as a physical consequence of the Sammis-Star trip, which itself was a correct functioning of the system as designed to the physical grid state prevailing at the time.

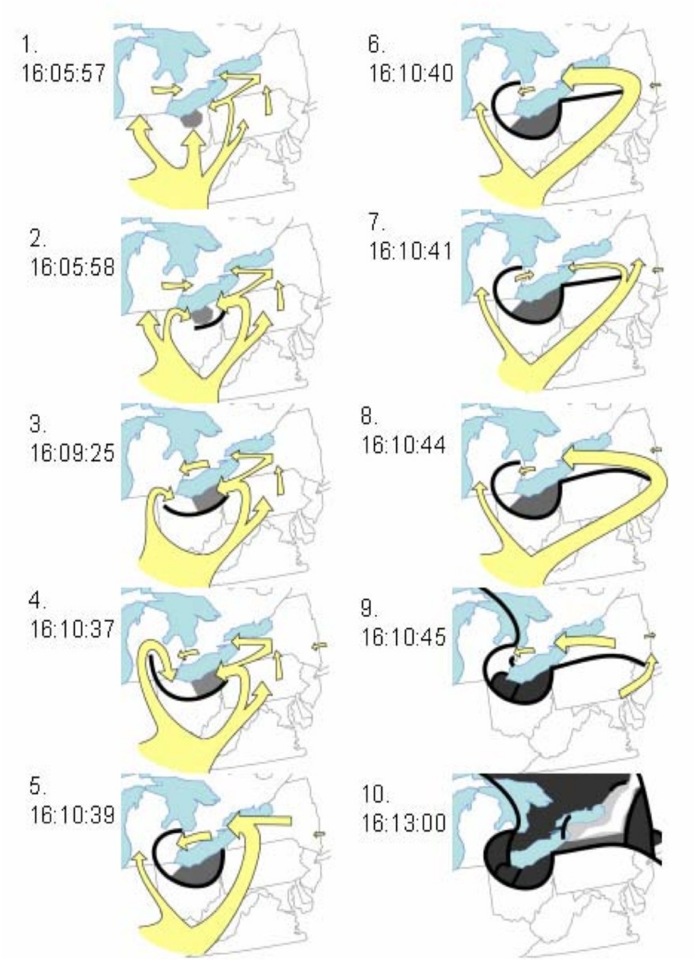

Figure 13: Power flows from immediately before the Sammis-Star trip (shown in diagram 1; Sammis-Star trip is in diagram 2) to complete system collapse (NAERC 2004, reproduced under fair use)

The trigger circumstances of normal accidents is often an emergent condition – it is rarely the case that such are known in advance, because, when they are, they would have been classified as a hazard and normal practice would have required mitigation or avoidance measures to be taken. During this incident, it appears also that the status of being a single point of failure under given grid conditions is also emergent.

.

2.2 History of the 2003 Event: Computer Problems Hindering Information Flow

It follows from the above conclusion in the report (NAERC 2004) that the most causally-significant events, namely those causal events where intervention was both possible and could have been effective in avoiding the blackout, occurred previous to the Sammis-Star trip. We recount the significant precursor events, illustrated in Figures 8 and 9 after the text.

- During the day, voltages were sagging in the Cleveland-Akron region and the system was judged retrospectively by NERC to have been approaching voltage collapse. However, this condition was not causal to the blackout
- Power transfers were "*high, but within studied limits and less than historical values*" (*op. cit.*, p12)
- At around 12:00, several lines in SE Indiana tripped
- At 13:31, Eastlake Unit 5, a generating station on the south shore of Lake Erie, tripped (Figure 14)
- At 14:02, the Stuart-Atlanta 345kV line tripped (Figure 14). This line was not in the MISO area, so MISO had no information on it, but its non-function caused the MISO state estimator to operate incorrectly.
- FE was during the entire time a major importer of power (*op. cit.*, p20). In the metropolitan area of south Lake Erie, air-conditioning loads were "consuming" reactive power, of which Northern Ohio was then a net importer. The system was not "reliable" with respect to reactive power, but this state was not causal in the blackout.
- At 14:14, the FE operators lost the "alarm function", on the computerised Energy Management System (EMS). The alarm function is an audible and visible annunciation of problematic status of some piece of kit (line, generator, capacitor bank, and so on). Operators remained unaware of this loss until the failure of the second EMS at 14:54, and did not realise that they had in fact lost alarm function 40 minutes earlier. There was no technical indication of loss of function. There had been calls from other operators, which hinted that the system state was not fully understood at FE, but these interactions apparently had "little effect". The EMS continued to exercise supervisory control and send correct status updates to other entities, including MISO and AEP.
- Although there had been partial losses of the alarm function before, this was the first time that total loss of function occurred.
- Between 14:20 and 14:25, various remote control terminals in substations ceased to function. This was noticed only at 14:36 through on-site inspection at a substation.

- At 14:27, the Star-South Canton 345kV line tripped, and then reclosed, at 54% nominal load (Figure 14). At this point, the FE operators had begun to lose situational awareness.
- At 14:41, the primary EMS server failed. The server function was taken over by a “hot stand-by”, but, because the alarm process was stalled, this transfer caused this backup system to fail at 14:54 (it is not explained how a stalled process caused the EMS server to fail).
- The failure of two EMS servers apparently caused the refresh rate on operators’ screens to slow to almost a minute, compared with the usual refresh rate of 1-3 seconds. There was a “warm reboot” at 15:08, but a warm reboot does/did not restart the alarm function. When FE operators became aware of the alarm-function problem at 15:42, another warm reboot was attempted between 15:46 and 15:59. Operators were not, however, aware that this action would not restart the alarm function.
- The MISO state estimator normally runs automatically every five minutes. A real-time contingency analysis is also performed, less frequently. The state estimator takes real-time telemetry data and constructs a “best-fit” power-flow model from that data. A contingency analysis is used to alert operators if the system is running “insecurely” [sic]. The state estimator sometimes may not resolve, if information is inaccurate, or may also report a high degree of error (presumably, an estimate). Both tools were said to have been under development and “*not fully mature*” (*op.cit.*, p36).
- At 12:15, the state estimator reported results “out of tolerance” (*op.cit.*, p36), due to a line in Indiana which had tripped but which the state estimator recorded as still in service. This information was updated manually; a correct update followed at 13:00 at which point the state estimator resolved acceptably. However, to troubleshoot the problem, the MISO operator had disabled the automatic five-minute state estimation regime. He left his position. The fact that the state estimator was not running regularly was discovered at 14:40. When the state estimator was rerun, it failed to resolve.
- The likely cause of the non-resolution of the MISO state estimator at 14:40 was the tripped Stuart-Atlanta 345kV line. This line is outside MISO’s area of responsibility and its status is not automatically linked to the MISO state estimator. There was a repeated failure to resolve until the MISO operator called PJM at 15:29 to determine the line’s status. After updated to the correct status (tripped), the MISO state estimator then resolved. Contingency analysis was run manually and resolved at 15:41.
- The MISO state estimator and contingency analysis were “*back under full automatic operation and solving effectively*” by 16:04. (*op.cit.*, p37). However, this was only a couple of minutes before the Sammis-Star trip and the start of the cascade.

- The MISO state estimator and contingency analysis were thus "*effectively out of service*" between 12:15 and 15:41 (*op.cit.*, p37). The report concludes, reasonably, that the lack of MISO diagnostic support contributed to the lack of situational awareness at FE.
- At 15:05, the Chamberlin-Harding 345kV line tripped and then locked out (Figure 15).
- At 15:32. the Hanna-Juniper 345kV line tripped and then locked out (Figure 15).
- At 15:41, the Star-South Canton 345kV line crossing the FE/AEP boundary tripped and locked out (Figure 15).
- The first two of these trips were not recognised by FE because of the loss of alarm function
- These trips obviously degraded the condition of the system.
- Between 15:39 and 16:08 there was a localised cascade of tripped 138kV lines in Northeastern Ohio
    - Seven lines tripped
    - Then the Dale-West Canton line, whose tripping caused the Sammis-Star 345kV line to overload, which initiated the blackout cascade irreversibly.
    - Then three more

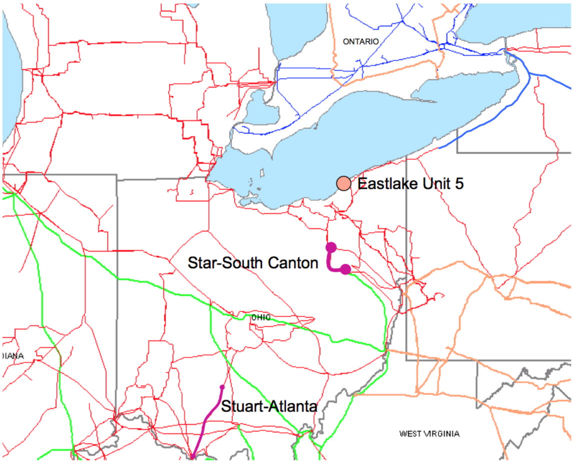

Figure 14: Initial line and plant trips (NAERC 2004, under fair use)

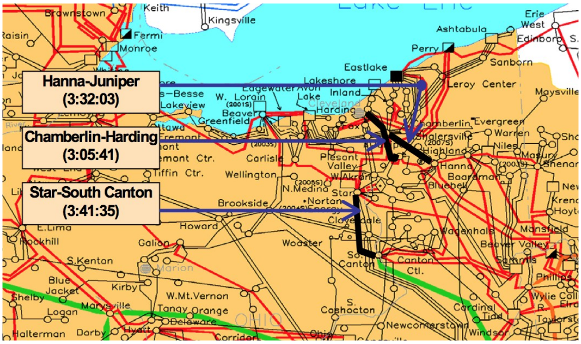

Figure 15: Three 345kV line trips (NERC 2004, under fair use)

The Sammis-Star trip at 16:05:57 EDT and the resulting alteration of network flows into the Cleveland-Akron metropolitan area are shown in Figures 10 and 11. This event is considered in the report to be key, in that the consequences could not have been avoided within the sociotechnical control system as it was at that time (*op.cit.*).

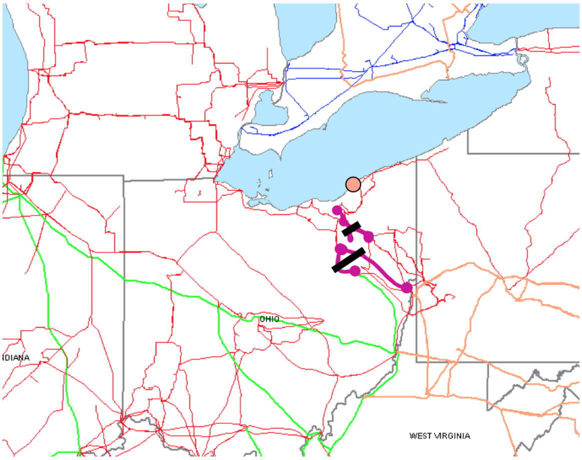

Figure 16: Cleveland and Akron supply cuts, through Sammis-Star Trip (NAERC 2004, under fair use)

The sequence of events after the Sammis-Star trip is detailed, and to our minds interesting, in particular how islands formed in sequence. A detailed set of maps showing the progression may be found in (*op.cit.*), also (Ladkin 2015), along with a synopsis of the full sequence of events from the NAERC report. We do not have space to consider it here. It contributes to our conclusions only through the observation of the rapidity of the decline to blackout. The entire sequence, covering multiple U.S: and Canadian states, from the Sammis-Star trip to the blackout in Figure 12, occurred within very few minutes of the Sammis-Star trip.

To a system analyst, the sequence of events before the Sammis-Star trip as related above are notable for their contrast with the causal factors as identified in the report. Effective human load management is key to the resilience of the system. Active management at times of high and variable load is key to load management. It is apparent that operators relied on the EMS for system management to the extent that loss of an alarm function significantly reduced the effectiveness of their management. It is also apparent that reliable state estimation is also key to effective management at times of high and variable load. The EMS, and the MISO state estimator, are computational systems whose functions are critical to resilience at times of high and variable load. However, it appears these systems were neither identified nor treated as mission-critical:

- Loss of an EMS function on which operators had ostensively come to rely, the alarm function, was not annunciated. Had the EMS system been identified as mission-critical and the usual design criteria applied, such a loss would have been annunciated.
- The MISO state estimator was temporally skewed for up to two and a half hours during the course of the incident, and did not reflect the true system state when skewed. This condition was likewise not annunciated, except indirectly, when the system failed to resolve, likely because a “hidden input” had a different value (the Stuart-Atlanta line). However, it was not annunciated over this two-plus hours that the estimate was static , a time when active management would have been required to ensure resilience.

It follows that the actual resilience of the Northeastern Interconnection is reliant upon the resilience of computerised systems, the FE EMS and the MISO state estimator and presumably equivalent kit at other operators and reliability organisations, which were not at the time either identified or managed as mission-critical. It is curious that this dependence was not addressed in the incident reports and analyses (*op. cit.*)

2.5 The November 2006 European Blackout

Except for a few DC lines, almost the entire European grid is a synchronised, in-phase AC grid. As in North America, the control stations of the network operators in Europe can usually see (part of) the states of the different components of the network, such as current load in megawatts (MW) and/or amperes (A), as well as the load limits for these lines. They will typically also be able to see how much is produced and consumed where, and also the states of the switches. High-voltage lines have automatic circuit breakers which will disconnect the line in case of overload. In contrast to the event we have just considered, in the European incident there do not appear to have been any computer anomalies causing misinformation. But not all the available information was used, and some important and faulty decisions were not checked.

At any point in time, the so-called *N-1 criterion* must hold in the European electric energy distribution grid. It means that "*any single loss of transmission or generation element should not jeopardize the secure operation of the interconnected network*" (BNA 2007). Closely related is the notion which we might call "*N-1 resilience*", that grid remains resilient when one element is lost. The term "*N-1*" comes from a single loss. N-2 resilient would be that the system remains resilient when two elements are lost. N-2 would be a much stronger criterion. The November 2006 incident showed that the grid at that point in time was N-1 resilient but not N-2 resilient.

All it took was a tall ship. On the evening of November 4 2006, a large cruise ship that had been built at the Meyer shipyards in Papenburg, on the River Ems near the north-east German coast, was scheduled to be conveyed along the river towards the North Sea.

There are several high-voltage lines passing across the River Ems, underneath which ships built in Papenburg have to pass on their way to the open sea. Most of the lines have been raised to allow safe passage, but for some ships the clearance of some lines still is not sufficient. Long-standing practice was to turn off some of these lines during a launch to allow the ship to pass. The disconnection had been requested for this ship some time in advance, and had been tentatively agreed by E.ON, the operator of the line in question. E.ON also informed the operators of neighboring network sections of the event.

On the day of passage, the shipyard requested shutdown of the line three hours earlier than originally planned. The earlier time was deemed more favourable and, after co-ordination with the neighboring operators RWE and TenneT, the request was granted by E.ON. RWE and TenneT checked for the fulfilment of the N-1 criterion prior to giving their approval to the disconnection, but E.ON did not do so (BNA 2007). In communications with RWE, it was found that, as a consequence of the disconnection, another high-voltage line, Landesbergen-Wehrendorf, which connected the networks operated by E.ON and RWE, was close to its load limit. In an attempt to reduce the load on that line, E.ON operators coupled busbars at a switching station. Because of the urgency of the situation, this action was not co-ordinated with RWE. Instead of reducing, the load on the line rose, and two seconds later the overloaded line tripped.

Because of the rapidly-changing distribution of the current flow, other lines in Germany and other parts of Europe tripped in quick succession, illustrated in Figure 17, which caused the transmission network to be split into three parts.

Nr.	Zeit	kV	Leitung
1	22:10:13	380	Wehrendorf-Landesbergen
2	22:10:15	220	Bielefeld/Ost-Spexard
3	22:10:19	380	Bechterdissen-Elsen
4	22:10:22	220	Paderborn/Süd-Bechterdissen/Gütersloh
5	22:10:22	380	Dipperz-Großkrotzenburg 1
6	22:10:25	380	Großkrotzenburg-Dipperz 2
7	22:10:27	380	Oberhaid-Grafenrheinfeld
8	22:10:27	380	Redwitz-Raitersaich
9	22:10:27	380	Redwitz-Oberhaid
10	22:10:27	380	Redwitz-Etzenricht
11	22:10:27	220	Würgau-Redwitz
12	22:10:27	380	Etzenricht-Schwandorf
13	22:10:27	220	Mechlenreuth-Schwandorf
14	22:10:27	380	Schwandorf-Pleinting

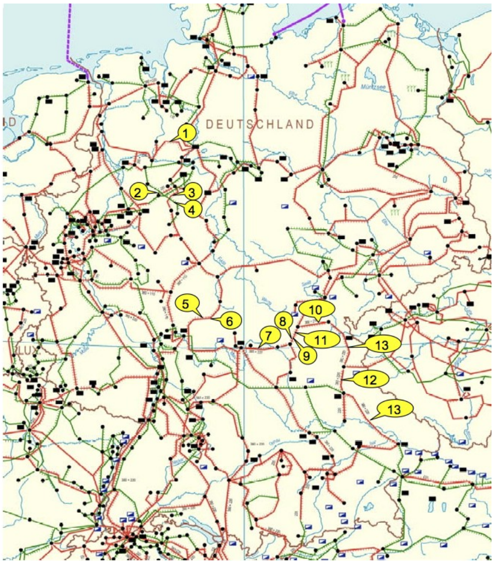

Figure 17: The first 14 lines tripped within 14 seconds (EON 2006)

Due to a discrepancy between production and consumption, frequencies in these areas began to drift apart, making a quick reconnection impossible. Generators were shut down and consumers had to be disconnected. More than 15 million people in Europe were affected by the blackout.

At more than one point, E.ON operators did not carry out a computer-assisted flow analysis before performing actions which altered the load distribution in their network. It appears that they relied instead on their experience to assess the state and security of the grid. The first instance was when the approval was given to disconnect the line over the River Ems. The second instance was when they decided to couple the busbars to alleviate the load on the Landesbergen-Wehrendorf line. However, even experienced operators cannot judge the behaviour of highly-complex interconnected systems intuitively.

A Why-Because Analysis of all the causal factors is available (Sieker 2008).

2.6 Total Power Blackout in the Swiss Federal Railways (SBB) Network on June 22, 2005

On June 22, 2005, the Swiss Federal Railways suffered a total power blackout. In contrast to the other two blackout events above, the initial events in the Swiss incident happened rapidly, within a few seconds, leaving operators almost no time for intellectual analysis but maybe just time enough to react. However, an aspect of system design leading to an “alarm flood”, with alarms required to be manually discharged before any action could be taken, contributed to the severity of the event.

In most electric-railway power grids, it is possible for the trains both to draw power for operation and to feed power back into the grid during braking. In order to avoid overloading the lines, the voltage of the line is measured, and a decision is made whether or not feeding power back in would be safe or not. In normal operation, this energy recuperation during braking both saves electrical power and reduces wear on the mechanical brake systems.

The sequence of events is elaborated in the report (SBB 2006). Two out of three power lines between two regions of the Swiss railway power network were shut down according to schedule due to construction work. The one remaining line tripped at 17:08h because of overload; there was no power connection between the Gotthard region and Central Switzerland. The railway power grid was separated into two islands, “North” and “South”. The South island was overproducing electricity, and an attempt at transferring power into the 50-Hz-network failed. Most generators were shut down automatically within seconds. All SBB railway operations in the canton of Ticino and at the Gotthard ceased.

In German-speaking and Western Switzerland, production in the powerplants Chatelard, Vernayaz and Etzel was increased. In concert with transfer from the Deutsche Bahn, the underproduction could be temporarily compensated. At 17:35, the coupling to the network of Deutsche Bahn was shut down. Remaining power stations in German-speaking and Western Switzerland further increased their power output, but ceased operations shortly after 18:00. Railway operations stopped in the North island as well. The islands are shown in Figure 18.

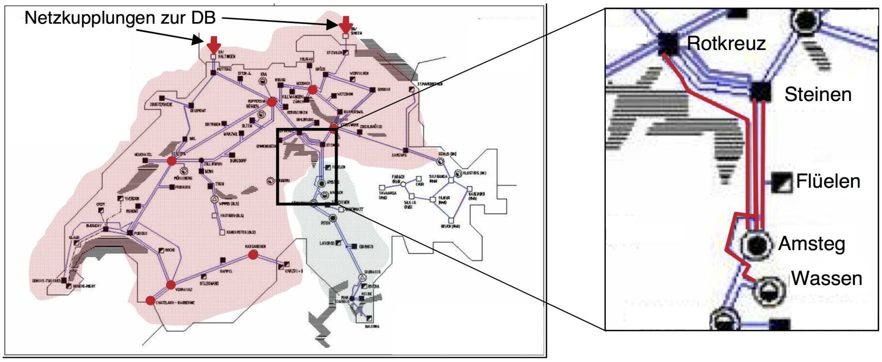

Figure 18: Grid islands which formed during the June 2005 SBB blackout (SBB 2005)

There were three main causes identified by SBB’s analysis (SBB 2006).

1. *Inappropriate risk estimate due to incorrect parameter values.*
   Wrong device parameters were a causal factor for an inaccurate risk analysis. The control centre assumed that the high-voltage line Amsteg-

Rotkreuz, which subsequently tripped, had a capacity of 240 MW. Although the line itself did in fact have a thermal capacity of 240 MW, the circuit breaker was set to 211.2 MW, limiting the usable capacity to this lower value. This latest current information was not available.

2. *Impossibility of timely and accurate assessment due to alarm flooding.*
   There were four individual alarm messages about the overload of the couplings to the network of Deutsche Bahn, but these were inundated under the flood of other alarm messages. In the first 60 minutes after the first line failed, 18,000 messages, including 3,400 critical messages accumulated in the control centre (SBB 2006). A filtering of messages was not possible, and each message had to be acknowledged manually, individually, before the status display of all network components in the control centre was updated. An early recognition of the alarms about overload of the couplings to DB would have allowed the timely reversal of the energy flow through the frequency converters from the civil 50-Hz energy grid to augment the missing power in the railway network. This reversal could have been completed in a few seconds. Instead, the transformers continued operating in "rigid" mode, supplying railway power into the 50-Hz grid.
3. *A Scenario like this had never been considered.*
   The possibility of a complete country-wide blackout of the railway power supply had never been considered prior to this incident, and was never included in operative risk management. Consequently, no contingency plans had been in place to prevent such an occurrence, or to minimize its consequences. Existing documentation about the prevention of (partial) blackouts proved unhelpful, because they were not tailored to the magnitude of this incident.

Of particular interest here is also the role of the N-1 criterion defined above. The N-1 criterion was knowingly disregarded, partly due to economic considerations (SBB 2006). When two of the three lines in Reusstal were shut down, the criterion was clearly violated, although continuing stable operation in both island networks would have been technically possible.

3 Conclusions Concerning Resilience

We wish to draw some straightforward conclusions. First, three observations.

1. In all three incidents, information was available to operators which, had it been acted upon, would have averted the blackout or mitigated its severity.

2. In two incidents, misleading information was displayed to operators and acted upon (unhelpful actions were taken; helpful actions were not taken). In one incident, some operators did not check available information but rather acted on an assumption which turned out not to hold.
3. In all three incidents, the generation and presentation of critical information was not subject to what critical-systems engineers would regard as appropriate assurance of dependability. In one incident, a design feature of the system inhibited timely action (the "alarm flood").

First, it follows from Observation 1 that the grid system considered as a physical system is theoretically resilient. All three incidents could have been avoided or mitigated through appropriate use of available information by operators. Second, considered as a sociotechnical system in which the actual behaviour rather than some idealised behaviour of human operators is taken into account, the system is manifestly less resilient than theory suggests. Third, the actual resilience of the sociotechnical systems could be significantly improved by routine critical-system engineering: identifying mission-critical functions in system components and ensuring their availability, or at least that their unavailability is not masked.

As things stand, the resilience properties of electricity grids are emergent. Considered as a physical system with ideal operator behaviour, say during design-time analysis, a grid appears adequately resilient. As an actual sociotechnical system, we have seen three cases in which it is in fact less resilient than supposed.

This situation in which presumed resilience is affected by actual implementation is also seen in other engineering domains. For example, air traffic control ground communications are effected though dedicated services. These services are often contracted out to telecommunications service providers, which run the dedicated services along with public telephone service and other services through non-dedicated equipment. A failure of this equipment, which may not be considered critical by the service provider, also causes the dedicated critical ATC services to fail (Neumann 1991). A contrasting case is that of motorway auffahr-accidents. The first author has shown, using Rational Cognitive Model checking, that auffahr-accidents are an emergent property of the motorway system-of-systems itself (Ladkin 2011).

We conclude from the three examples we have considered that the actual resilience of some sociotechnical systems is lower than a theoretical analysis might have led engineers to believe. We have also observed that conditions which causally lead to a normal accident are often emergent. Finally, we have observed that being a single point of failure under certain conditions is often an emergent property. It follows that the resilience properties of these systems are in large part emergent.

Acknowledgements A talk based on a version of part of this paper, as well as events from motorways and finance, was given at an SCSC seminar in London in April 2015. The first author thanks Michael Parsons for the invitation to present there, which gave the impetus to prepare a preliminary version of this paper. Both authors thank him for his helpful comments on the first version of this paper.

References

Berlekamp E R, Conway J H, Guy R K, (1982) Winning Ways for Your Mathematical Plays, Academic Press, 2nd edition A K Peters, 2001-4.

Bundesnetzagentur (2007) Report by the Federal Network Agency for Electricity, Gas, Telecommunications, Post and Railways on the disturbance in the German and European power system on the 4th of November 2006. Technical report, Bundesnetzagentur

Conway J H (1970) Life (game). Explained in (Berlekamp et. al. 1982). Also see https://en.wikipedia.org/wiki/Conway%27s_Game_of_Life , accessed 2015-11-14.

E.ON Netz GmbH (2006) Bericht über den Stand der Untersuchungen zu Hergang und Ursachen der Störung des kontinentaleuropäischen Stromnetzes am Samstag, 4. November 2006 nach 22:10 Uhr. Technical report, E.ON Netz GmbH.

Ladkin P B (2001) Causal System Analysis, e-book available from http://www.rvs.uni-bielefeld.de/publications/books/CausalSystemAnalysis/index.html , RVS Group, University of Bielefeld

Ladkin P B (2011) The Assurance of Cyber-Physical Systems: Auffahr Accidents and Rational Cognitive Model Checking, RVS Group, University of Bielefeld, December 2011. Available from http://www.rvs.uni-bielefeld.de/publications/Papers/20111230CPSV2.pdf , accessed 2015-11-15.

Ladkin P B (2015) Synopsis of the 2003 North American Blackout, available from http://www.rvs.uni-bielefeld.de/publications/Reports/LadkinSynopsis2003Blackout.pdf , RVS Group, University of Bielefeld

Meyer J (2009) Defining and Evaluating Resilience: A Performability Perspective, slides from a talk at PMCSS9, September 17, 2009. Available from http://web.eecs.umich.edu/people/jfm/PMCCS-9_Slides.pdf , accessed 2015-11-15.

U.S. National Transportation Safety Board (2014) Auxiliary Power Unit Battery Fire, Japan Airlines Boeing 787-8, JA829J Boston, Massachusetts January 7, 2013, Report AIR-14-01, NTSB 2014.

Neumann P G (1991) AT&T phone failure downs three New York airports for four hours, Risks Forum 12.36, September 1991. Accessable from http://catless.ncl.ac.uk/Risks/12.36.html#subj1.1 accessed 2015-11-19.

North American Electric Reliability Council (2004) Technical Analysis of the August 14, 2003, Blackout: What Happened, Why, and What Did We Learn?, Report to the NERC Board of Trustees by the NERC Steering Group, July 13, 2004. Available from http://www.nerc.com/docs/docs/blackout/NERC_Final_Blackout_Report_07_13_04.pdf)

Perrow C (1984) Normal Accidents: Living with High-Risk Technologies, Basic Books, 1984. Updated edition, Princeton University Press, 1999.

Schweizerische Bundesbahnen (SBB) (2005) Strompanne der SBB vom 22. Juni 2005. Technical report, Schweizerische Bundesbahnen

Schweizerische Bundesbahnen (SBB) Zentralbereich Revision (2006) Second Opinion zur Strompanne der SBB. Technical report, Schweizerische Bundesbahnen

Sieker B M (2008) European Electricity Blackout, November 2006. Causalis Technical Report, 2008. Available from http://www.causalis.com/90-publications/EuropeanElectricityBlackout.pdf , accessed 2015-11-15.

U.S.-Canada Power System Outage Task Force (2004) Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations