Normal Accidents and Computer Systems

Michael Dorner
Advisor: Heiko Niedermayer
Seminar Future Internet WS2013/14
Chair for Network Architectures and Services
Department of Computer Science, Technical University Munich
Email: dorner@in.tum.de

ABSTRACT

Accidents happen to all of us. Whether it be small or big ones, sometimes there seems to be no way to avoid them. The more complex our systems get, the harder it gets to understand what caused an accident, and how it could have been prevented. And yet, sometimes the risk of failure is so high, that accidents must not happen. Normal accident theory and high reliability organization are two scientific approaches to deal with accidents in high risk systems in order to predict and explain, respectively prevent accidents. However, both theories come from an era, where analogous technology was dominating, so its its unclear whether they still hold in a world dominated by digital technology. By exploring the implications of the use of computers in traditional systems as well as the implications of HRO and NAT for fully digital systems, it will be shown, that both theories can be applied in both cases with slight modifications and restrictions.

Keywords

accidents, high reliability theory, normal accident theory, critical infrastructure, Internet

1. INTRODUCTION

We live in a world full of amazing technology, which allows us to impact our environment in ways, that would have been considered nothing but magic only a century ago. While most technology was developed to improve the quality of our lives, failure in one of our new technologies may also result in a disaster. The potential extent of such an accidental catastrophe has increased proportionally with the sophistication of our systems and now includes the possibility of total annihilation of life as it is on this planet, trough accidental nuclear war. In order to better understand and hopefully prevent possible accidents caused by new human technology science has made an attempt to better understand these accidents. While accidental nuclear war is actually a concern in this area, it surely also is an extreme case to stress the importance of this research. Normal accident theory (NAT)[13] and high reliability organization (HRO)[15] are two of the more prominent pieces of scientific work which deal with system reliability, organization and accidents. Both originated in the investigation of classical high-risk systems and system accidents in chemical reactors, flight control or nuclear power plant (npp) security. Based on empirical studies, they try to better understand what causes accidents in complex systems respectively what needs to be done to avoid them. However, both are theories based on accidents and systems from a time when a lot of devices were analogous and computers often played only a minor role. Since the processes within computer systems exist within a virtual reality which we created as part of our reality, it is not natural to assume that both theories apply to them in the same way they apply to classical engineering. Thus the goal of this work is to investigate which impact both theories have on our digital systems.

Before heading into this discussion, there will be an introduction into both theories, as well as a short summary of a number of accidents, which were used to motivate both theories. Since the main contributors to both theories had a very spirited discussion about the relation of HRO and NAT, it is necessary to take a look at their relation and find a position suitable for applying both theories to modern computers and computer networks. Subsequently the applicability of both theories in digital systems will be discussed along with the possible consequences one should draw from their application. The paper is finalized by a summary of the findings.

2. BASICS

Neither NAT nor HRO can be considered common knowledge, so it is only normal to explain both shortly before discussing them.

2.1 Normal Accident Theory

Normal Accident Theory originates from Charles Perrow's investigation of an accident at Three Miles Island(TMI), an American nuclear power plant (NPP), which underwent a partial meltdown in 1979. Although it did not cause a larger nuclear disaster like Chernobyl or Fukushima, the accident resulted in an extensive investigation. In the light of this investigation Perrow voiced the opinion[12], that in tightly coupled systems with highly complex interaction accidents through unexpected interaction are normal, thereby giving birth to "Normal Accident Theory". In his subsequently published book on "Normal Accidents" Perrow did not only lay out his theory, but also defined a framework for

doi: 10.2313/NET-2014-03-1_04

the classification of accidents. While a full summary of his framework would be far too extensive, it cannot be avoided to cover some definitions, which will matter later on, when comparing NAT and HRO. In his definitions Perrow divides systems into units, parts, subsystems, and the system itself. He introduces a separation between "component failure accidents", which involve "one or more component failures, that are linked in an anticipated sequence"[13], and "system accidents", which "involve unanticipated interaction of failures"[13]. The system accidents are the ones Perrow considers "normal accidents", since they cannot be prevented according to his theory. It is important to mention that despite their name, "normal accidents" are an exceptional, rare kind of accidents according to Perrow. Another important aspect of his theory is the definition of tight coupling and complex interaction. Perrow defines complex interaction through:

- *Local proximity*: Components are close to each other, thus causing failures in one to possibly affect the other simply because of their proximity.
- *Common-mode connections*: Mainly characterized by their impact on a number of otherwise independent systems. A common-mode failure is the single source for failure in multiple systems, although they may not seem to have a connection with each other at first glance.
- *Interconnected subsystems*: Subsystems are connected with each other, thus likely propagating failure on the subsystem level.
- *Limited substitution of materials*: Materials cannot simply be replaced, making leakage or breakage problematic.
- *Unknown/unfamiliar feedback-loops*: Feedback loops are based on the idea, that the system receives feedback and adjusts its behavior. In the best case this change in behavior is to produce output, which is closer to or is the desired output. However, having a feedback loop within a system, which is not planned for or which is not effectively blocked, if it is not desired, can cause the system behavior to deviate from the expected behavior.
- *Multiple and interacting controls*: Control is e.g. performed via multiple terminals, thus requiring coordination.
- *Indirect information sources*: It is not possible to directly deduce the actual system state from merely observing it, but indicators have to be monitored.
- *Limited understanding of processes*: The process does not allow complete understanding due to either a large number of possibilities or non-deterministic elements.

Among these criteria, especially common-mode connections and feedback loops will frequently be present, but not perceived in complex systems. Systems, which are not complex, i.e. which do not have these characteristics are referred to as "linear" by Perrow.

Just as for complexity Perrow also defined criteria for tight coupling:

- Processing cannot be delayed
- Fixed Order of Sequences
- Only a single method leads to success
- Little slack in resources
- Buffers and redundancy have to be present by design
- Substitution of resources has to be designed in

These characteristics should mostly be self-explanatory, so it is not necessary to give an extensive explanation at this point. One thing, that may not immediately be clear however, is the notion of designed-in buffers/substitutes: components in tightly coupled systems are not easily replaced, and buffering is not possible without components, which are explicitly included by design to allow for it - think about an empty bucket to catch liquid from a leaky pipe and try to find something that does the same for an overheating nuclear fuel rod.
Perrow's work gained a lot of attention back when it was released, and constituted a new point of view on accidents, which was picked up by the a group of researchers in Berkeley in their work on HRO[15]. It was this HRO group's findings, which Scott Sagan later contrasted against Perrows original work in [18]. While the discussion his paper started will be analyzed later on, only Sagan's contributions to NAT are of interest at this point. Perrow himself acknowledged and praised two major contributions to his work by Sagan[14]: the fact, that he outlined a difference in the theoretical models HRO and NAT had used, and his emphasis on group interest as a relevant factor. While Perrow merely clarifies that he was implicitly assuming a *nonrational* garbage can model (see Appendix A) as underlying organizational model in response to the first contribution, he points out Sagan's second finding as a novel, interest-theoretical aspect of his theory. The inclusion of group interest means that some groups, which are part of the governing process of an organization, may put other things first instead of safety (e.g. profit). Although the work cited above is typically a decade or two old NAT is still commonly referenced in scientific articles[17] as well as news articles about major accidents such as Fukushima [11] or financial crisis[20], thus keeping it relevant.

2.2 High Reliability Organization

Since accidents are the main topic of this paper, one may wonder, why it is necessary to introduce a second theory at this point, when NAT is an accident theory of its own. Depending on which position one takes, it is either, because HRO is not a theory of its own, but a complementary organizational strategy to provide high reliability in systems prone to normal accidents, or it is because HRO is a competing theory, which suggests that normal accidents can be prevented by sticking to certain guidelines (in that case one would likely call it HRT - T for theory). However, before it is possible to take either side, it is necessary to understand what HRO is about. High Reliability Organization

was initially driven by two groups of researchers: one at the University of California, Berkeley and the other at the University of Michigan. The "Berkeley group's" main members were LaPorte, Rochelin and Roberts; the "Michigan group" was mainly represented by Weick and Sutcliffe. The research of both groups, which frequently reference each others research, is centered around ways to establish a safety culture in organizations, which operate systems, which are prone to normal accidents. From their own observations of aircraft carriers and flight control[15, 2] they identified a set of properties, which, according to them, help those organizations improve their reliability, thus making them "High Reliability Organizations". Interestingly there is no real agreement on a definition of reliability among the HRO researchers, and it seems that they currently have a rough consensus on reliability as "the ability to maintain and execute error-free operation"[17]. The original research of the Berkeley group[15] and their subsequent work[3, 2, 4, 6] finds the following requirements for an organization to become a HRO:

- Redundancy: Especially the early HRO research outlines the importance of redundancy in personnel and safety mechanisms to better cope with component failures and also to ensure that decision-making involves more than a single operator.
- Prioritization of safety: Government and organization leaders put reliability and safety first. Other aspects, including performance may suffer, but leaders accept the loss of performance to achieve higher reliability and safety.
- Organizational Learning: The organization learns from ongoing operation, thus continuously improving the ability to deal with failures.

The Berkeley group had become a little quiet over the last years, and there is no recent publication from them. Opposed to that the Michigan group still seems to be working actively on their HRO-model, which is fairly frequently updated (last 2011), in their book on "Managing the Unexpected". The focus of the Michigan group is more on the investigation of organization culture, which HROs have to establish. It currently lists five important aspects of organizational culture, which HROs should embrace[21]:

- *Preoccupation with failure*: Members of the organization are not focused on what confirms their ways, but what opposes it, i.e. possibilities to fail are actively perceived, people pay attention to possible new or unknown modes of failure and learning takes place. It reduces overconfidence and encourages a state of mindful operation.
- *Reluctance to simplify*: Organization members are animated to stay wary of the complexity of the system they operate, and thus make more considerate decisions, although likely lowering performance. Some processes, which could in theory be done by a single person may e.g. be performed by a group of equals without a shared perspective to profit from their collective understanding.
- *Sensitivity to operations*: Members are aware of the current situation and its implications.
- *Commitment to resilience*: The organization's capabilities to improvise and react to new situations are constantly maintained and improved.
- *Deference to experience*: The most experienced members make decisions despite hierarchies if failure happens. Decision making is decentralized but based on the culture put in place by centralized organization leaders.

These criteria for an organizational culture are not solely the work of the Michigan group, and is difficult to track which group outlined the importance of a certain aspect at first, but they seem to agree in large parts that these qualities are important to establish a culture in which an organization can operate at high reliability. The main difference between both groups is, as already mentioned, the increased focus on organization culture by the Michigan group, whereas the Berkeley group was also still involving some systems design considerations such as redundancy. In general it is important to outline the relevance of *decentralized, rational decision making* for HROs and their focus on reliability and safety over budgets and performance. LaPorte from the Berkeley group notes in some of his post-Cold-War work on reliability-oriented organizations, that "when either the consensus about their value declines or economic resources in general become more dear, reliability regimes are more difficult to sustain, especially after conspicuous success and/or as system resources become relatively more scarce."[6]. HROs are faced with the challenge to maintain their reliability record under these conditions.

2.3 Important Accidents

After explaining NAT and HRO, we will now shortly take a look at some accidents, which have influenced them and which have drawn attention to this kind of research. The literature on accidents and reliable organization lists a large number of other accidents, and there have also been more fatal ones than those, that are about to be explained, but they make good examples to show up important aspects of both theories, and have thus been chosen. Because most accidents are covered extensively in accident reports created by experts, we will stick to a short description. We will thus take a look at the following points:

- What was the starting point?
- What happened?
- What were the causes for failure?
- What were the consequences of the accident?

2.3.1 Three Miles Island - TMI

What was the starting point: Three Miles Island was and is a npp in Pennsylvania, US. It consists of two reactors, TMI-1 and TMI-2.
What happened: On the 28th of March 1979 there was a partial meltdown in TMI-2 and a small amount of radioactive gas was released

What was the were the causes for failure: Due to previous maintenance work, all feed-water pumps (primary, secondary, emergency), which are required to transport coolant to the reactor were offline. Thus no heat was deduced from the reactor the pressure rose, since pressure and temperature are proportional by the laws of physics if the volume remains the same. An automatic valve opened to reduce the pressure inside the reactor, and should have closed after pressure returned to normal, but instead was stuck open due to mechanical failure. Naturally opening the valve primarily reduced the volume of the coolant inside the reactor, thus leaving it open leads to a lack of coolant. The plant operators failed to recognize this situation, for one because of a misleading indicator light, which displayed wrong information (valve closed), and also because they were preoccupied with the correctness of this light and consequently ignored indicators, which should have let them realize the true nature of the failure they were facing. Ultimately, the lack of cooling caused the nuclear fuel rods to overheat and ultimately the partial meltdown. A detailed description of this accident can be found in [13, 10] and in many other sources, since this accident was very well investigated and a lot of information is available to the public.
What were the consequences of the accident: Luckily, the damages to the environment, as well as the exposure of the population to radioactive material were low, such that it is commonly agreed upon today, that the TMI accident had no observable long-term consequences for the health of the surrounding people. As already mentioned there was a government investigation, which also resulted in Perrow's basic paper on normal accidents. Anti-nuclear protests gained credibility from this accident, especially in the US, but the consequences for the nuclear industry were insignificant. TMI-1 is still operating today, and has a license, which lasts at least until 2034. TMI-2's decontamination officially ended in 1993, although it is still monitored.

2.3.2 *The Bhopal Disaster*

What was the starting point: A pesticide factory in Bhopal, India, surrounded by slums.
What happened: Large amounts of highly toxic gas leaked into the air.
What were the causes for failure: Unlike TMI, the Bhopal accident is not fully resolved until today. What is for sure is that water somehow entered a tank full of methyl isocyanate (MIC), a deadly gas. Water and MIC cause an exothermic reaction, which results in increased pressure. The high pressure lead to the release of several tons of MIC into the air through emergency relief valves. Government investigation found that leaky pipes and valves were most likely to be the reason, and that water had gotten into one of the tanks while they were flushed for cleaning. It further indicates, that the plant was in a horrible condition with several safety measures being non-functional, employees completely untrained to react to accidents, and the plant being understaffed. Cost appeared to be the driving force behind these shortfalls. Although frequently mentioned for its tragic outcome, the accident itself seems to be poorly researched, most likely due to the lack of an independent investigation. The company operating the plant originally claimed that sabotage must have been the reason - very likely to avoid compensation claims - but the government's assessment is considered a fact today.
What were the consequences of the accident: This accident is the often referred to as "worst industrial disaster" in history, and has caused at least 3,787 deaths according to the local government, although others claim, that the number of deaths caused by it are around 25,000. The number of injured people is estimated to be around 500,000 - 600,000 and the area is still not decontaminated. Without taking sides on the cause, it is clear that the surrounding environment of the plant, where many people from the plant lived in slums, the failure to inform surrounding inhabitants of the gas leak, and the lack of an evacuation strategy lead to the disastrous outcome of this accident. As already mentioned severe negligence was very likely the reason for this disaster. An Indian court, in agreement with this point of view, found eight former plant employees guilty of "death by negligence"[19] in 2010 and sentenced them to two years prison and a fine of 2,000$.

2.3.3 *Challenger*

What was the starting point: The Challenger was one of the Space Shuttles of the US Space Program **What happened**: The Shuttle was torn apart by the aerodynamic forces mid-air after its launch on the 26th of January, 1986
What were the causes for failure: Two redundant o-rings, did both not seal one of a tank correctly due to cold weather, gas leaked, and one of the rockets used to boost the shuttle during take-off was no longer correctly attached to the space-vehicle. The resulting changes in aerodynamics increased the physical forces to an extent, which exceeded the limit the shuttle could take. Problems with the o-ring were known to the manufacturer beforehand, but instead of grounding all space shuttles, they added this behavior to the acceptable conditions, because it would work under normal conditions. **What were the consequences of the accident**: Subsequent flights were canceled and all shuttles were grounded for 32 months. A commission was mandated to investigate the accident. The commission found the design of the o-ring to be faulty and thus NASA or the manufacturer should have grounded all shuttles until the issue was resolved. On a side note: Richard Feynman, a member of the commission, was so appalled by NASA's "reliability culture"[1], that he insisted on adding personal notes to the report, for it to have his name on it. The report [16] is still publicly available from NASA.

3. NORMAL ACCIDENTS AND HIGH RELIABILITY ORGANIZATIONS - A CONTRADICTION!?

During the explanation of both HRO and NAT, it was already mentioned that the authors of both theories had an argument about whether HRO complemented NAT, or whether it was a competing theory of its own. The discussion between Perrow and La Porte was sparked by Sagan's work on the "Limits of Safety"[18]. While the main topic of his book was the safety of nuclear weapons and defense systems, he had also investigated both HRO (referred to as HRT by him) and NAT in this context, and found that they are competing theories, which exclude each other. Perrow agreed with Sagan's analysis and complimented him on his work and his contributions to NAT. LaPorte and Rochelin on the other hand, representing the Berkeley group, did not agree at all

and explicitly addressed Sagan's as well as Perrow's arguments in a paper with the sole purpose to contradict them. Of all the points made in this discussion there are three, which seem to touch central issues of both theories the most, and which will be investigated to justify the position, which will be taken with respect to both theories for the rest of this paper.

3.1 Possibility of Error-Free Operation

The possibility of error-free operation is something that obviously contradicts NAT, which claims that some accidents cannot be avoided. Therefore if HRO would actually claim to offer a way of error-free operation, this would already be the point where we could stop and side with NAT, because HRO's claims would be implausible for the present and impossible to prove for the future. As a matter of fact, it is true that La Porte and Rochelin have claimed to have learned "the degree and character of effort necessary to overcome the inherent limitation of securing consistent, failure free operations"[5], but La Porte claimed that Sagan had misunderstood this statement. According to his statement on the issue in [5] they thought that the required effort would be too big to be surmounted. Therefore, and also because HRO in general contains elements like learning from errors and near-misses which contradict the idea that error-free operation can be achieved, it should not be seen as realistic goal of HRO - otherwise it would likely also be referred to as total or complete reliability theory by its authors. Even Perrow seemed to think that both theories agree that error-free operation is not a realistic goalPerrowLoS. Interpreting the explanation of La Porte[5], the relationship between effort and gain in reliability assumed by HRO appears to be similar to the acceleration to the speed of light: the closer one comes to 100%, the more the effort required to come any closer increases. Both, tight coupling and highly complex interactivity, seem to increase the required effort even further. In connection with Perrow's explanations on how interactive complexity and tight coupling cause unpredictable interaction of failures in [13] it appears that both experienced the same phenomenon, but HRO focused on how organizations dealt with this challenge, while Perrow payed more attention to what factors contribute to it. Ultimately, it is safe to say that both HRO and NAT agree, that error-free operation is not possible.

3.2 Effectiveness of HRO Methods

Another thing that NAT theorists had criticized about HRO, was the effectiveness of their methods in general. Sagan and Perrow both voiced the opinion that some of the techniques HRO relies on do not have any provable beneficial effect. While some of their criticism directly addresses concepts of HRO, they also doubt the benefits of their methods in general.

3.2.1 Specific Criticism

First of all, we will turn to Perrow's criticism of specific methods which HRO suggests. His specific criticism in [14] addresses thee concepts:

Centralization and Decentralization. This point is one of the few where HRO and NAT truly conflict, and where it is not possible to convincingly argue for either side. While HRO suggests a centrally imposed HRO-culture which is executed in a decentralized fashion by all organization members, NAT expects that both models cannot be combined, Perrow explained[14]. While he considered both concepts essential to deal with complex interactivity respectively tight coupling, he also thought that they cannot be combined. As already mentioned is possible to side with either party here, but HRO's model offers more opportunities, as it considers mixed forms of both concepts a possibility and also dynamic shifts from one concept to the other. Since La Porte el al give credible proof that this can work in reality, e.g. aircraft carrier operation, their opinion is just as justifiable as Perrow's, who refers to his theoretical explanation in[13]. HRO's more dynamic approach to decision making, which is also offers the possibility to shift decision making in centralized organizations, e.g. from the highest ranking to the most experienced person, as the principle "deference to experience" dictates it, seems more promising than simply surrendering to the fact that aspects of two mutually exclusive concepts are required.

Training. While training for emergencies is intuitively a necessary measure, the implications of HRO's understanding of training go way beyond regular emergency drills: it expects that organizations forgo routine and stability in exchange for challenge and variety to improve the experience of the employees with irregular circumstances. While this may make sense in some situations, Perrow's argument[14], that this is not an option for systems with especially high risk like npps, is a striking one for high risk organizations. Even in regular organizations this is unlikely to be an option because it will likely decrease productivity. It is however not surprising that La Porte et al observe this kind of behavior in organizations like an aircraft carrier, which practically does nothing but training in times of peace. For regular organizations the kind of learning HRO suggests does seem unreasonable to implement though. While it is certainly not wrong to stay wary and have the preoccupation with failure HRO-culture demands, intentionally mixing up regular operation seems just unreasonable for most organizations.

Learning. The aspect of organizational learning is closely related to that of training, since more training would obviously result in more learning. Therefor if learning was an effective measure, the relevance of training would also increase. Perrow also criticizes this aspect of HRO based on an extensive list of examples where organizational learning did not happen[14], and with an earlier study on accident investigations, which found that accident-investigations typically only investigate those sub-systems which failed and not the role of other sub-systems in this failure, which may obstruct learning. The latter argument is also supported by the fact that accident investigations often stop after assigning blame - often to the operator - as [7] found. This hunt for a scapegoat which follows many accidents is also a reason why the full set of failures and their interaction will likely remain undiscovered. Therefor a solely beneficial effect from this kind of organizational learning cannot be ascertained in general. Learning from biased investigations may even worsen the error handling of the organization. The arguable benefits of learning from accidents and near-misses also further limits the benefits one should expect from the kind of training HRO suggests. Therefor siding with NAT on this arguments seems to be the better choice.

doi: 10.2313/NET-2014-03-1_04

3.2.2 General Criticism

General Applicability of HRO Methods. What Perrow and Sagan criticized most in the HRO-methods is that they come from organizations which have not experienced failure. HRO tried to determine factors, which allow them to achieve this high reliability; Perrow refused this approach as "selecting on the dependent variable"[14]. He thought that just because the organizations observed have shown common approaches during failure-free operation, this does not mean that those approaches are helpful to achieve failure-free operation in general. The fact that a future accident may expose a new, previously unpredictable cause, i.e. it is a normal accident, is what makes NAT practically impossible to falsify and any error-preventive measure impossible to prove. Proving error-free operation measures seems to be an undecidable problem, because it would require the prove of future properties of a system, which is typically undecidable, although we will forgo a proof of this property here. In consequence it is true, that HRO's usefulness cannot be proven beyond the point of no doubt, but it is very common to apply theories, which show desirable effects in practice until they are proven wrong. Examples are the application of mathematical theories in finance and politics, like game theory, large parts of all social sciences, which are not proven conclusively in a manner satisfactory for many STEM-scientists, and the different models that have been used to describe atoms throughout the 20th century, which were wrong or incomplete and yet allowed for major scientific advancements. It is also the fact that severe negligence of HRO principles has caused some of the most serious accidents, that supports HRO's claim to enhance reliability: among the three examples listed earlier on, both the Challenger and the Bhopal disaster could have probably been prevented by a more HRO-influenced organization. Because HRO methods could have likely prevented or reduced the extent of these accidents, it would be unjust to generally dismiss them, because they come from organizations which have not failed, especially could have prevented accidents.

Applicability of HRO Methods to Normal Accident-Prone Systems. The second general criticism of NAT's advocates is that HRO is not applicable to those systems, where "normal accidents" are especially likely to occur. According to Perrow these systems all have a high degree of coupling and interactive complexity, which the organizations HRO investigated have not. Perrow and La Porte rant on about this classification issue quite a while in [14] respectively [5]. As explained earlier, these two criteria are determined by a set of factors, which Perrow nicely outlined in his book[13]. The problem is that "high" and "tight" are not defined in any objective way, therefor their meaning is completely bound to the subjective perception of the person looking at an organization. It is like asking two people to name a big number: if you ask a computer scientist he may come up with something like 2^{128}, while a normal person might just say one million. Perrow's comments on this topic in his book show, that he is fully aware of the subjectivity of his categorizations and the fact that he is lacking a metric[13]. Given these considerations, it is very likely, that the HRO groups, having a different mindset than Perrow, simply came to a different categorization. In fact it seems that nothing but the subjective estimation of the person applying NAT serves as the function, which projects organizations into this fourfold table. Analogous to the principle that ambiguities in a contract are held against the party who put up a contract in US law practice, the subjective definition of when a system is tightly coupled with highly complex interactions should be held against Perrow, therefore voiding his argument that HRO does not investigate organizations with tight coupling and highly interactive complexity, simply because they are not well enough defined. HRO's applicability to what Perrow considers "normal accident-prone" is still limited, due to the effort required to further increase the reliability of systems, which are already very close to error-free operation. Since Perrow himself declares that some of the recent history's worst accidents like Bhopal, the Challenger-crash, the Exxon Valdez and Chernobyl were not normal accidents[14], HRO would be a great contribution if accidents like those could be avoided through it.

4. ACCIDENTS IN COMPUTERS AND COMPUTER NETWORKS

After an extensive discussion of two major pieces of work on accidents/accident prevention the actual matter of accidents in computer networks can now be addressed. The two theories presented are not the only ones of interest to this area, but their relationship is complex enough already and the introduction of other models such as Leveson's STAMP[9, 7] or even her more basic work on "Safeware"[8] would simply add to the confusion. All of the previously presented systems have huge catastrophic potential for the real world, something which is sometimes said to be a difference between traditional systems and computer networks. Computers and computer networks are sometimes thought of as a world of their own, although they have already begun to have a notable impact on real world objects. The fact that we have seen not any major accident caused by computers does not mean that they do not carry the potential for causing such accidents. With smart grids and the Internet of Things coming some of our everyday life is already moving into the digital world, but technologies which are a bit further away like autonomous cars or robots will definitely mean that digital systems can cause just as catastrophic accidents as analogous ones. For accident theory, computer technology introduces two different aspects: computers which conquer the domain of classical systems, and computer systems themselves, which follow entirely new laws with respect to their internal operation.

4.1 Computers in Classical System Accidents

Classical systems are the ones that rely on clever use of physics and other sciences exploring the laws of nature, e.g. chemistry to make our natural system behave in a the way we want it to. Because nature is not exactly obedient, those are also the systems that carry higher risk the more extreme our nature hacks are. For these systems the introduction of computer systems often means that control is taken from a human and given to a computer, which is merely supervised by a human to ensure its correct operation. This change in control is often going to add complexity, because the human supervisor gets less direct information from the system and if the computer is networked and uses information provided by other computers this may introduce new ways of propagating error throughout the system. But in general a computer is a component that can fail just as much as any

other part. Most operators do no have complete knowledge of the system they supervise at every level, but know the in- and outputs, and possibly intermediate results their system will produce. Because these results are not physical in the case of a computer, communicating the exact situation to the outside will be an important challenge. If the internal state is not communicated correctly and clearly to the outside, computers will significantly impair HRO-operation, because operators might not notice a near-miss and may not be able to stay aware of what is happening inside the system as HRO demands it. Furthermore the much faster development lifecycle which often changes existing systems drastically may reduce the benefits of experience. Yet one should not expect that the "culture of reliability" looses importance, because there will still be people operating the system somewhere in the background which need to pay even closer attention to whether the system runs as expected. We can expect that the effort will increase though, since we have added yet another level of complexity to our systems. While HRO faces some challenges in its transition to the digital world, NAT is golden. Because one of its assumptions has always been that adding new means to prevent accidents to a system will also introduce new ways of failure, it is not to be expected that it looses any ground. Since computers are no more than a part of traditional systems, their internal behavior is not as important at this point as the behavior they show to the outside, i.e. the signals they send and receive to and from other components. Since we are not just switching to digital for fun, but because of the huge potential of this technical innovation, we should also expect to see new modes of failure, especially when we try to handle failures that analogous systems could experience through "smart" systems. NAT assures us that even when we go smart, we will see failure that those smart devices wont be able to handle.

4.2 Accidents in Computer-Systems

4.2.1 Software and System Accidents

Software accidents, are accidents where software behaves in an unspecified way. Most often they are bugs, and thus simple component failures. Software bugs are quite difficult to combine with HRO, since HRO relies on a certain mindset which people inside HROs share, which is not applicable to computers as they are not. Therefor HRO does not apply to software and computer systems. Software and the internal workings of a computer also limit NAT to some extent, because "normal accidents" are caused mainly by things that are unexpected or not well understood, which may not exist in a system that we created ourselves. The only reason why we could expect "normal accidents" in software on a single computer is non-deterministic execution which causes data-races, but if enough attention is payed to synchronize this should not happen. Besides, many of the factors that increase complexity or tight coupling are factors that computer design limits by design e.g. through out-of-order execution, best effort service, and dynamic scheduling. Therefor the probability for normal accidents on a software level is quite low, although it may well increase when we build more advanced systems, where timing guarantees and similar features matter. In single, non-networked computers, normal accidents should therefor be possible to prevent. Networked computers are very likely different: if we take the Internet of today for example we see a fairly loosely coupled system, which has best effort service and optional reliability in transport. As soon we use TCP to gain reliable transport, we add a certain level of complexity to the Internet, because all of a sudden we have a feedback loop, namely the TCP congestion control mechanism, which complicates things significantly, as it influences the shape of the traffic. If we add more guarantees to the Internet such as QoS-guarantees, we complicate the interaction of the packets even more and if they go over the same link with dynamic bandwidth allocations, the Internet may well become complex enough for us to loose oversight and experience "normal accidents". That should not stop us from advancing our technology, but we should stay conscious of the fact that complicated systems will show failure that we cannot anticipate. What complicates this situation even more is that we cannot effectively monitor the Internet with respect to how data flows globally and therefor our feedback from the Internet may be poor or incomplete. Because the Internet is the backbone of many of the amazing visions for the future, anything that is built on top of it, has the base level of complexity and coupling that comes from the Internet. All of the technologies built on top add to this base level may worsen the situation, especially cloud computing, which is probably one of the best examples for a common-mode connection. While cloud-based data-warehousing is a huge trend at the moment, mission-critical services should thus not run at a single data-center to prevent catastrophic common-mode failures. If the service is important enough, e.g. power supply, using a dedicated physically separate network may even be a good choice to make sure the complexity from the Internet cannot cause unexpected failure. For the global Internet it is impossible to say whether normal accidents will or will not become a major issue, because it is hard to predict what the Internet will evolve into. Although it looks like managed services and QoS will gain importance, it would be a common interest of all users to keep the basic network as simple as possible to keep the base complexity and coupling low.

4.2.2 Computer Operation and Administration

Apart from flawed software or systems themselves, there is also the aspect of their operation. While software may behave as specified if it is patched correctly, errors may arise if it is misconfigured or not administrated correctly. HRO's principles are generally applicable to the operation of computer systems as much as they are applicable to any other organization, with the exception that they require communication between operators and administrators in IT. While administration must put an emphasis on offering a stable environment for its users, they should pay attention to bugs and issue that may indicate systematic problems. Although this combination should usually allow administration to improve the experience of its operators, this synergy is often limited by the fact that operation and administration are not done by the same organization. The trend towards outsourcing administration or buying things -as-a-Service in the cloud make HRO almost impossible to realize in many scenarios, because the administration has extremely limited information about the system state form the user point of of view, which usually offers more insight into the weaknesses and problems of the current system. Since the communication between users and administration is typically limited to complaints, this makes it very hard for administration to realize the culture of reliability. Although HRO is a bit hamstrung by this, this area should usually not be very prone

to "normal accidents" because most of our systems consists of many small and independent machines, which are quite linear.

5. CONCLUSION

In the end there is much to learn form both NAT and HRO. While NAT mainly highlights the importance of paying attention to the level of complexity and coupling we introduce by our designs, it also reminds us to stay wary of the possibility of unexpected failures. HRO teaches us the importance of decentralization and the focus on reliability. Its culture of reliability should prove valuable to almost any companies in the future in operation and is expected to maintain its effectiveness in traditional systems even though it may require extra effort to compensate for the culture-free decision making of computer systems. The combination of both theories during design and operation can surely help organizations on- and off-line avoid accidents, although they seem to leave some room for other theories especially when it comes to design which avoids normal-accidents and operation strategies, which do not rely strongly on human intelligence.

6. REFERENCES

[1] Feynman R.Appendix F - Personal observations on the reliability of the Shuttle, in "Report of the PRESIDENTIAL COMMISSION on the Space Shuttle Challenger Accident"

[2] LaPorte, Todd R. :*The United States air traffic control system: increasing reliability in the midst of rapid growth*, 1988.

[3] LaPorte T.R. : *High Reliability Organizations: The Research Challenge*, HRO Project Paper, Institute of Governmental Studies, University of California

[4] La Porte, Todd R. and Consolini, Paula: *Working in Practice But Not in Theory: Theoretical Challenges of High-Reliability Organizations*, in Journal of Public Admin- istration Research and Theory, 1, 1991, pp. 19-47

[5] La Porte, Todd R. and Rochlin, Gene I.: *A Rejoinder to Perrow*, in Journal of Contingencies and Crisis Management Vol. 2, Nr. 4, 12/1994

[6] La Porte, Todd R..: *High Reliability Organizations: Unlikely, Demanding and At Risk*, in Journal of Contingencies and Crisis Management Vol. 4, Nr. 2, 06/1996

[7] Leveson, Nancy :*A new accident model for engineering safer systems*, in Safety Science 42.4 pp.237-270, 2004

[8] Leveson, Nancy :*Safeware: system safety and computers*. ACM, 1995.

[9] Marais, Karen and Dulac, Nicolas and Nancy Leveson:*Beyond Normal Accidents and High Reliability Organizations: The Need for an Alternative Approach to Safety in Complex Systems*, Engineering Systems Division Symposium, MIT, Cambridge, MA March. 2004

[10] Nuclear Research Commission(NRC): Backgrounder on the Three Miles Island Accident, http://www.nrc.gov/reading-rm/doc-collections/fact-sheets/3mile-isle.html

[11] Osnos, E. : *The Fallout*, The New Yorker, 11/17/2011, http://www.newyorker.com/reporting/2011/10/17/111017fa_fact_osnos?currentPage=all

[12] Perrow, Charles.: *The President's Commission and the Normal Accident*, in D. Sils, C. Wolf and V. Shelanski, Accident at Three Mile Island: The Human Dimensions, Westview, Boulder, pp.173-184

[13] Perrow, Charles.: *Normal Accidents*, Princeton University Press, 1999

[14] Perrow, Charles: *The limits of safety: the enhancement of a theory of accidents., in* Journal of contingencies and crisis management 2.4 (1994): 212-220.

[15] Rochlin, Gene I., Todd R. La Porte, and Karlene H. Roberts : *The self-designing high-reliability organization: Aircraft carrier flight operations at sea*, in Naval War College Review 40.4: 76-90, 1987

[16] Rogers Commission:*Report of the PRESIDENTIAL COMMISSION on the Space Shuttle Challenger Accident*,, http://history.nasa.gov/rogersrep/v1ch4.htm

[17] Rouncefield M. and Bubsy J.: *D 7.2.1*,2013

[18] Sagan, Scott.: *Limits of Safety: Organizations, Accidents, and Nuclear Weapons*, Princeton University Press, 1993

[19] Soutik Biswas:*Bhopal trial: Eight convicted over India gas disaster*, BBC News, http://news.bbc.co.uk/2/hi/south_asia/8725140.stm

[20] Surowiecki, J. : *Bonds Unbound*, The New Yorker, 02/11/2008, http://www.newyorker.com/talk/financial/2008/02/11/080211ta_talk_surowiecki

[21] Weick, K., and Sutcliffe K.: *Managing the unexpected: Resilient performance in an age of uncertainty* Vol. 8. John Wiley & Sons, 2011.

Appendix A - Garbage Can Model

The garbage can model is a model for organizational decision making, which was originally established by M. Cohen, J. March and J. Olsen. It has been adapted later on, and thus multiple versions can be found. The characteristic trait of this model is however the fact that it **assumes decision making to be based on stochastic events** involving a set of streams, such as policies, politics and rather than rational analysis.