Big Data for Reliability Engineering:
Threat and Opportunity

Vitali Volovoi
Independent Consultant
vitali@volovoi.com

Abstract - The confluence of several technologies promises stormy waters ahead for reliability engineering. News reports are full of buzzwords relevant to the future of the field—Big Data, the Internet of Things, predictive and prescriptive analytics—the sexier sisters of reliability engineering, both exciting and threatening. Can we reliability engineers join the party and suddenly become popular (and better paid), or are we at risk of being superseded and driven into obsolescence? This article argues that "big-picture" thinking, which is at the core of the concept of the System of Systems, is key for a bright future for reliability engineering.

Keywords - System of Systems, complex systems, Big Data, Internet of Things, industrial internet, predictive analytics, prescriptive analytics

1. INTRODUCTION (SO CLOSE, SO FAR AWAY)

It is not easy to be a reliability engineer: his (or much less likely, her) place does not neatly fit into the basic architectural design of the "house" of engineering. The main building blocks of this house are long, vertical, highly specialized pillars (that can involve a particular discipline, such as structures, or a particular module, such as a component), combined with the "roof" of systems engineering.

The roof deals with the "big picture" and aims to combine the pieces of the puzzle provided by the individual disciplinary pillars. Systems engineering is positioned at the interface of management and engineering, and as a result enjoys a somewhat privileged position. It is recognized as an important competitive advantage for technology companies, so the latest trends are highly publicized and studied in business schools. Indeed, one may recall the "Six Sigma" of yesteryear, or marvel at the modern hype related to Big Data analytics and the Internet of Things (IoT).

From the technical perspective, reliability engineering falls under the general umbrella of systems engineering (that includes overlapping fields such as industrial engineering, decision and management science, operational research, or more recently, analytics). It shares with the rest of the fields under this umbrella the need to abstract away most domain-specific information, and to use tools that are mainly domain-independent[1]. As a result, it increasingly shares the *lingua franca* of modern systems engineering—probability and statistics that are required to balance the otherwise orderly and deterministic engineering world.

And yet, reliability engineering does not wear the fancy clothes of its sisters. There is nothing privileged about it. It is rarely studied in engineering schools, and it is definitely not studied in business schools! Instead, it is perceived as a necessary evil (especially if the reliability issues in question are safety-related). The community of reliability engineers consists of engineers from other fields who were mainly trained on the job (instead of receiving formal degrees in the field). This community is quite conservative; the field can hardly be described as trend-setting. Even if fashionable business trends directly affect reliability engineering (e.g., total quality management, or Six Sigma), the impact from the reliability engineering perspective is mostly reactive.

There are multiple causes for this status quo, and the study of those causes would be the worthy subject of a separate article. Briefly, those causes are discussed next.

2. HOW DID WE GET HERE?

Modern reliability engineering can be traced back to the 1950s, spurred by the complexity of emerging computer systems and military equipment [1]. In the social sciences there is an interesting notion of nested natural cycles; in particular, the Kondratieff cycle is related to the lifespan of a generation

---

[1] Here the operative word is "mainly." For example, in selecting the parametric failure distribution, distinguishing between underlying physical mechanisms can be critical (e.g., the weakest-link nature of a failure can point toward a Weibull distribution in the case of cracks, while corrosion spreading proportionally to the size of the existing damage could be modeled naturally using a lognormal distribution). The distinction can easily be lost given the proliferation of curve-fitting software (see the discussion of the resource curse later in the article).

of technological innovation. If a single biological generation of the work force spans about 30 years, the Kondratieff cycle spans two biological generations, or about 60 years [2].

The first half of the cycle is characterized by rapid innovation in and dissemination of the new technology, while the second half is characterized by the maturation of the technology and acquiring incumbent status. Our society as a whole is deemed to be in the middle of the Kondratieff cycle associated with the information technology cycle beginning in the middle of the 1970s [2]. However, it can be useful to look at the engineering world's own Kondratieff cycle related to operations research that started a generation earlier [3]. It is difficult to overestimate the importance of technological innovation driven in large part by Cold War competition as well as the uniquely beneficial economic position of the United States after World War II. We will also note that the IEEE was founded in 1963.

Two complementary forces largely shaped reliability engineering as we know it today: the separation of incentives between acquisition and sustainment (the principal-agent problem), and "small data" capacity constraint. Both of these effects are described in more detail next. (Spoiler alert: the relevance of both forces is greatly reduced, if not eliminated, in the current environment.)

A. Principal-Agent Problem

Let us consider a complex engineering system, such as a gas-fired power plant, or a train. The direct upfront cost of acquiring such a system usually constitutes about 30-40% of the total life-cycle cost, with the balance (60-70%) stemming from operation and support (O&S) costs[2]. However, the seller—such as the original equipment manufacturer (OEM) —usually has little or no incentive to reduce O&S costs. In fact, spare parts and repairs can be a significant source of revenue and profits for the OEM. This is commonly referred to in economics as the principal-agent problem. The customer (the principal) is subject to information asymmetry: he or she is more interested in knowing what O&S costs are going to be, while actually having less information about those costs compared to the agent (the OEM). For example, the OEM might not share reliability data with the customer. Complex systems are sufficiently unique (in terms of their configuration and operation profiles) to obscure the information signal to a large degree.

Over the past 30 years or so, there has been a clear trend toward the customer preferring, and often demanding, the transfer of the risk associated with uncertain O&S costs to the OEM. On the other hand, businesses often perceive it to be advantageous for them to climb up the food chain and provide a service rather than a product (e.g., cloud computing); as a

result, long-term service agreements became the norm for complex engineering products, thus aligning the incentives for improving reliability for OEMs or their equivalents.

B. Small Data Constraint

In his famous article "More is Different" [5], P.W. Anderson discusses the pitfalls of a reductionist approach to dealing with complex systems, and the need for a hierarchical view with loosely coupled layers. The sheer number of entities that comprise a given layer, along with the complexity of their interactions, effectively preclude a "constructivist" bottom-up approach, and instead require a separate set of laws (rules) that govern the "emergent" behavior. It is therefore fruitless, for example, to construct biology laws using chemistry, even though individual entities do react in accordance with chemistry's laws.

Unlike natural complex systems, the principles of the operation of engineering systems are relatively well understood and documented, so long as a system operates as intended by its designers. As a result, for a single system there is a clearly defined hierarchy of "indentures" (that can have several layers) all the way down to an elementary part (such as a bolt), each carrying a specific functionality. If the operating conditions of an elementary part are known, it is relatively easy to understand the reliability of that part in isolation. This part can be subject to accelerated tests, and the data requirements for establishing such reliability are relatively modest. Standards and guidelines (such as MIL-HDBK-217 and IEC-TR-62380) can be established and even maintained. If the reliability of parts is known and the interactions among the parts follow simple (linear) rules, the design life of a system can be inferred.

How well this designed reliability matches operational reliability is altogether a different question, as here a well-ordered (mostly linear[3]) engineered world meets reality. This abstraction is, however, a necessary simplification, given the limitations of data collection and processing. In the past, attempts to look at more complex models of interactions among system components largely failed, due to a lack of data that could support those models and describe interactions more precisely.

---

[2] For example, estimates for yearly U.S. defense O&S spending are about $150 billion, as compared to a $100 billion annual budget for procurement of new equipment [4].

---

[3] Here linearity refers to a functional relationship among the system components from the reliability perspective. Of course, that does not preclude other non-linear relationships. For example, an air-conditioning system can have a temperature controller that maintains a desired temperature by using feedback, thus implementing a nonlinear relationship. One can refer to this as an inner-loop control. However, from the reliability perspective, one can abstract away the nonlinear subtlety of the operation, and focus on binary working/nonworking possible states of the controller: the outer loop does not have feedback.

However, one can imagine a different world, where the stress and temperature for that bolt are actually measured and recorded as a function of time throughout the operation of each system where such bolts are installed. One can also imagine that each bolt is constantly assessed in terms of its wear and presence of cracks. It is clear that in this alternative world the analysis tools would be different.

In the case of a bolt, this alternative world is not a reality yet (and it might not be financially prudent to make such a world a reality even if it were technically feasible). However, for a large number of components in modern engineering systems, the historical footprint of some environmental (usage) conditions, as well as internal "health" conditions, can indeed be collected and stored. Furthermore, the resulting terabytes of data can be effectively processed using modern Big Data processes that enable effective parallelization, such as MapReduce and Spark.

The implications of such condition-based maintenance are not always straightforward, especially in the current "transitional" period where operational strategies often outpace corresponding reliability analysis. Paradoxically, in some critical systems, the introduction of condition monitoring can weaken the incentives for understanding reliability issues. Indeed, consider a safety-critical item in a traditional engineering system. The designed reliability of this item has a direct impact on the safety of the system, and therefore a significant effort is made to evaluate its reliability: if the reliability of the item is less than designed, the safety risks increase. This situation can be contrasted with a modern system that provides condition-monitoring capabilities for the same safety-critical part. Instead of a periodic inspection of the part, a warning signal is sent to a maintainer indicating that the part needs a replacement. So, the safety of the system is ensured, but reliability can suffer, because the safety and reliability issues get decoupled, and therefore there is less incentive to make the part reliable.

In general, the trend is toward the increasing sophistication of operational technology (OT), providing a flexible and automated means of reconfiguring a system in order to provide uninterrupted service to the customer. As OT strategies become smarter, reliability assessment must keep up and take into account more sophisticated "outer loop" control strategies. In this context, using the reliability methods of the past is often insufficient for assessing the reliability of modern complex systems that act more and more like the complex natural systems discussed in the beginning of this section.

In summary, traditional reliability engineering managed to defy the messy nature of complex systems by confining uncertainty to component-level behavior and adhering to simple outer-loop control logic. But, quoting a classic, we are not in Kansas anymore.

3. BIG DATA IS WATCHING YOU

It is natural to take the hype generated by Silicon Valley with a grain of salt. Yet a closer look at the hype surrounding Big Data clearly reveals that "there is some 'there' there."[4] A confluence of several technologies that allow collecting, storing, and processing data at the scale that was unthinkable not so long ago is quite remarkable. To date, the applications that are most noticeable (and not necessarily in a good way!) are in the consumer area (such as custom-tailored on-line advertisements and flagging unusual credit-card activities).

However, both large engineering companies and major IT companies take notice and invest in what effectively amounts to a long-term commitment to building infrastructure for the Internet of Things (IoT) in general, and in the context of complex systems, for the Industrial Internet (or, as General Electric refers to it, the "Internet of really important things") [8]. Old-school companies like Lockheed Martin are busy hiring freshly minted computer science graduates as data scientists, who are often blissfully unaware of their engineering ignorance.

The long vision of the IoT promises a futuristic view that includes an automatically adjusted thermostat designed to help alleviate the peak electricity load on a hot day based on user-specified preferences. That vision might be realized some day, or instead we might use aesthetically pleasing home batteries courtesy of a (competing) vision form Elon Musk, that will smooth the demand in a different way. What is important, however, is that in the short term one of the most prominently featured use cases for IoT is a cost-effective improvement in the reliability of engineering systems [9].

So, if you are a reliability engineer, not only is Big Data watching you, it might also be after your job. It is hard not to notice how formerly independent companies (e.g., Relex, ReliaSoft) that provided reliability tools are being bought up by companies that aim at providing a "cradle-to-grave" digital tracking of engineering systems.

Of course, as we were told by politicians of all stripes, including John F. Kennedy, Al Gore, and Condoleezza Rice, in Chinese the symbol for crisis is the same as the symbol for opportunity[5]. And the community of reliability of engineers is indeed capable of correcting itself. For example, nowadays, Bayesian statistics gets a bit more respect than it used to [11].

4. SYSTEM OF SYSTEMS (SOS) TO THE RESCUE

It can (and should) be argued that reliability engineering was ahead of the curve in terms of providing a structured and balanced view of an uncertain world. On the one hand,

---

[4] The original expression "there is no 'there' there," indicating a lack of substance, is from Gertrude Stein's description of Oakland in the 1930s in comparison with San Francisco [6]. The expression was also used by Admiral Gehman, the chair of the Space Shuttle Columbia accident investigation board, to describe the lack of technical expertise at NASA in safety-critical areas [7].

[5] Not really [8].

deterministic logic was used to describe the interrelationships among system components that followed the designers' intent; on the other hand, the "irreducible" (also called aleatory) uncertainty was confined to the risk that some of the individual components might fail.

Such a divide-and-conquer strategy worked well for relatively simple engineering systems that were explicitly constructed in modular fashion. As engineering systems become more complex, their structures become more similar to those of natural systems. Importantly, there is compelling evidence that nature favors modularity as well [12]. More specifically, both natural and engineering complex systems are likely to be "nearly decomposable" into individual modules or layers. The coupling is "weak" in the sense of the rate of information exchange, but it is of critical importance, and it most certainly involves feedback mechanisms.

A recent large-scale quantitative study was aimed at understanding the nature of the successful forecasting of complex events [13]. Remarkably, one of the critical paths toward good forecasting was constructing an informal structured model, what was referred to as "fermi-izing" the problem (after Enrico Fermi, who advocated the method among his students)[14]. To this end, a European effort in so-called "process mining" is of great interest [15]. In contrast to data mining that deals with "black box" modeling (i.e., lacking any recognized internal structure), process mining aims at automatic learning of the model's structure appropriate for the modeled process (based on the historical "traces" of event sequences).

In other words, the ecosystem of reliability engineering is changing, and it behooves the reliability engineering community to adapt to the new challenges. In particular, a better understanding of "weak" couplings and their efficient modeling is of great importance to reliability engineering in general, and at the very crux of the SoS problem.

Another related important aspect of SoS is the need to keep the "big picture" in mind and maintain the correct level of abstraction. As computational tools become increasingly more powerful, it is always tempting to create large detailed models. This can be called a "resource curse"[6] as unfortunately, such models are of very limited general value.

Keeping the big picture in mind and avoiding getting stuck in silos is another quality of reliability engineering that is of universal and increasing value (see for example [16]), and again SoS issues fit right in. In fact, it is plausible that big-picture skills are the least likely to be replaced by a computer.

To illustrate the fundamental nature of the changes to the reliability engineering ecosystem, let us consider perhaps the most familiar emergent behavior of molecules in a volume of gas: if the system is closed, a stable thermodynamic equilibrium is reached, and a macro property (temperature) statistically represents the average molecular velocities (i.e., individual microproperties). The temperature provides a useful characterization of the behavior of a very large number of entities: it is a Dow Jones for molecules.

Classical engineering systems are designed to be as stable as a simple volume of gas. Instability is "designed out," and

---

[6] The term is originally used for countries blessed with natural resources that actually hold the countries back in their development.

reliability effectively measures the level of success in avoiding the instability. In contrast, complex systems must balance some degree of stability with the ability to adapt to changing environment. Here is an analogy of a medium for a signal: If the medium for a signal is too stable, the signal decays; if it is too unstable, the signal amplifies and grows too big. To be effective in propagating the signal, the medium should obey the Goldilocks principle—it must be just right. The resulting balance is called "equilibrium away from equilibrium," "the edge of chaos," or more technically, self-organized criticality (criticality being a technical term for this "edge" or boundary between stable/static and unstable/chaotic) or SOC [17]. SOC is a universal feature of complex systems.

These notions from theoretical physics seem esoteric and irrelevant to reliability engineers. In fact, they are of critical importance in general, and particularly so in the context of SoS. Next, some of these connections are briefly discussed.

As complex system controls become more sophisticated, a system increases its ability to maintain equilibrium at the desired level. When the limits of the system's control are finally exceeded, the loss of equilibrium is sudden and drastic. For example, let us consider a computer network that has a primitive routing algorithm. When one of the nodes (or links) fails, performance (measured, for example, by packet delays) degrades. If more nodes fail, the performance will continue to gradually degrade. In physics this is referred to as a second-order phase transition.

Let us now instead consider a "smarter" routing algorithm that is very efficient in rerouting the packets around the failed node or link. The end effect is that for an outside observer there will be no perceptible degradation of performance until the system is saturated, and there is a sudden onset of drastic congestion. In physics this corresponds to first-order transition. A classical example of such "masking" of the onset of the problem by means of a sophisticated control is putting out small fires, and therefore increasing the chances for a large fire. Needless to say, this masking can be potentially dangerous, as it can create a false sense of security.

Yet another intriguing example of this phenomenon is the role high-frequency trading (HFT) can potentially play in suppressing trading volatility under normal circumstances while facilitating large jumps during seemingly innocuous small disturbances (as occurred during the "Flash Crash" in 2010) [18].

Understanding and predicting such "extreme" events is a critical issue for SoS, and there is some encouraging research in this area. The prevailing wisdom about complex systems is that SOC dictates a power law distribution for event sizes. In log-log scale the frequency of events plotted against their magnitude is a straight line. Earthquake magnitudes, city sizes (Zipf's law), and internet connections all follow this distribution.

These power laws are also called scale-free due to the absence of the characteristic scale of the phenomena. In other words, extreme events are not any different than smaller events, except in just being bigger. This leads to the "black swan" hypothesis that extreme events are fundamentally unpredictable. An opposing view contends that there are certain important "pockets of predictability" where the emergence of extreme events is driven by global (on the system scale) synchronization [19]. As a result, these extreme events can possess distinct signatures that can be recognized

in advance. Further research on this system-wide risk phenomenon in the engineering context (rather than theoretical physics context) would be quite welcome indeed.

5. CONCLUSION

Technologies associated with Big Data bring fundamental changes to reliability engineering. Reliability engineers should stop worrying and learn to love Big Data. The reliability for SoS lies at the core of the most exciting and timely challenges that face the reliability community today. Addressing these challenges requires maintaining a big-picture view and avoiding getting stuck in silos. Then the new technological capabilities will allow more successful tackling of end-to-end metrics (e.g., delivering Systems of Systems functionality), and planning for undesirable emergent behavior. However, there are cultural and generational barriers that need to be overcome to make the renewal of reliability engineering successful.

REFERENCES

[1] J.H. Saleh and K. Marais, "Highlights from the early (and pre-) history of reliability engineering," Reliability Engineering and System Safety, vol. 91, pp. 249-256, 2006.

[2] T. Devezas and G. Modelski, "Power law behavior and world system evolution: a millennial learning process," Technological Forecasting & Social Change, vol. 70, pp. 819-859, 2003.

[3] W. Thomas, "History of OR: useful history of operations research," ORMS Today, vol. 42, n. 3, June 2015.

[4] S.I. Ervin, "Defense industry targets $150B weapons maintenance market," National Defense, July 2012.

[5] P.W. Anderson, "More is different," Nature, vol. 177, n. 4047, pp. 393-396, 1972.

[6] G. Stein, Everybody's Autobiography. New York: Random House, 1937.

[7] Space Shuttle Columbia Investigation, Hearing before the Committee on Commerce, Science and Transportation. United States Senate, S. Hrg. 108-838, May 2003.

[8] H. Kodesh, "The cloud and the internet of (really important) things" https://youtu.be/g3UX_WHbGcA

[9] http://www.itworldcanada.com/slideshow/10-internet-of-t hings-use-cases-that-have-already-become-cliches

[10] http://www.pinyin.info/chinese/crisis html

[11] R.A. Evans, "Bayes is for the birds," IEEE Transactions on Reliability, vol. R-38, p. 401, 1989.

[12] H.A. Simon, "Near decomposability and the speed of evolution," Industrial and Corporate Change, vol. 11, n.3. pp, 587-599, 2002.

[13] P.E. Tetlock and D. Gardner, Superforecasting: The Art and Science of Prediction. New York: Crown Publishers, 2015.

[14] P. Morrison, "Letters to the editor," American Journal of Physics, vol. 31, n. 8, pp. 626-627, August 1963.

[15] W. Van der Aalst, Process Mining: Discovery, Conformance and Enhancement of Business Processes. Heidelberg: Springer, 2011.

[16] G. Tett, The Silo Effect: The Peril of Expertise and the Promise of Breaking Down Barriers, New York: Simon & Shuster, 2015.

[17] P. Bak, How Nature Works: The Science of Self-Organized Criticality. New York: Copernicus Press, 1996

[18] A. Kirilenko, A.S. Kyle, M. Samadi, and T. Tuzun, "The flash crash: the impact of high frequency trading on an electronic market," SSRN, abstract 1686004, 2015.

[19] D. Sornette and G. Ouillon, "Dragon-kings: mechanisms, statistical methods and empirical evidence," Eur. Phys. J. Special Topics, vol. 205, pp. 1-26, 2012.

AUTHOR BIOGRAPHY

Vitali Volovoi is an independent consultant in the area of the industrial internet, risk, reliability, and the dynamic interactions of complex systems. He collaborates with Mitek Analytics, consults for Logistics Management Institute, and serves as a member of the NASA Statistical Engineering team. He has led many projects, including turbine analysis for the Air Force, quantitative risk assessment of the Space Shuttle wiring for NASA, air transportation safety analysis for NASA, condition-based maintenance of gas turbines for Siemens, and reliability of the Hubble Space Telescope gyros for NASA. He has received two NASA Engineering and Safety Group Achievement Awards, a NASA Engineering and Safety Technical Excellence Award in 2014, and the best tutorial award at the Reliability and Maintenance Symposium (RAMS) in 2011. He has a Ph.D. in Aerospace Engineering from Georgia Tech, and a University Diploma in Mechanics and Mathematics from Moscow State University.