

The survivability of survivability

S Dietrich, P Y A Ryan

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213
{spock,pryan}@cert.org

6 September 2001

"It's too bad she won't live. But then again, who does?" Gaff in Ridley Scott's "Bladerunner", 1982.

Introduction

Survivability is typically defined along the following lines: the ability to continue to fulfill a mission even in the face of attacks and failures. Crucial here is the acknowledgement that it is impossible to foil all attacks and prevent all failures. No single component of a system is immune to failure or subversion.

There appear to be a number of technical obstructions to progress towards the goal of designing and evaluating survivable systems:

Lack of precise definitions and focus

The above definition is fine as far as it goes, but to evaluate a system we need more precise, formal definitions. Some attempts have been made to provide at least frameworks in which such definitions could be formulated but these seem not to have been widely accepted.

Our first problem then is that we have no clear criteria against which to measure success or progress.

A process algebraic formulation using a weakened form of non-interference in which acceptable performance in the face of certain envelopes of attack/failures can be defined is proposed in [7]. This is presented in the process algebra CSP. Given the unbounded, dynamic nature of many of the systems of interest it may be necessary to use a dynamic calculus such as the pi-calculus or ambient-calculus.

Survivability is a relative attribute

It is clear that survivability cannot be viewed as a binary, absolute property. To boldly assert that a system is survivable is clearly absurd. A system can at best be survivable in the face of certain classes of attack or failure. Safes are traditionally classified according to how long they can be expected to resist certain types of attack, such as break-ins or fire. We need analogous measures for survivable systems.

One problem here is perhaps a psychological one: there is an ingrained tendency for people with an information security background to think in absolute terms. Much of the early work on information security sought absolute definitions of say secrecy. More recently there has been a recognition that absolute security is unattainable but the mindset lives on.

There seem to be two ways forward: to construct measures of acceptable performance in the face of certain classes of attack, as suggested above, or to construct survivability orderings, perhaps analogous to the security ordering of Jacob [5]. The latter is attractive mathematically but is perhaps not satisfactory from a practical point of view: simply knowing that one system is "better" than another is not really very satisfying.

A framework in which survivability orderings can be defined is presented in [8]. In essence a system P is deemed at least as survivable as Q with respect to some mission defined by M , written $P \geq_M Q$, if:

$$\forall \phi \in \Psi \ M \parallel P \parallel \phi \text{ refines}_{FD} \ M \parallel Q \parallel \phi$$

In other words, in the presence of any hostile environment ϕ from some class Ψ , and when restricted to behaviors associated with the mission M , P provided at least as much functionality as Q .

Traditionally the problem has been viewed principally from the defender's point. A healthier approach would be to take a more game-theoretic approach and view the problem from both the defensive and offensive perspectives. This is akin to the evaluation of a cryptographic device: you really need to adopt the cryptanalyst's mindset. This also acknowledges the relative nature of the attribute: there will always be an attacker with sufficient resources, time and determination to defeat any given set of countermeasures.

A key ingredient of survivability seems to be containment of damage: bounded attack effort should result in bounded damage. Containment should be in space and time: only a finite portion of the system or its functionality should be affected and the system should recover within finite time of attack activity ceasing. Note that damage might also be contained with respect to levels of abstraction of the system: damage manifest at one level might be masked at a higher level.

Perhaps a "universal" characteristic of survivability is the way damage scales with intensity of attack. Ideally damage should scale at most roughly linearly with attack effort. This seems to give a sort of general, dimensionless characterization. Coming up with natural measures of attack effort and damage can be a little delicate though.

Probability assumptions

An obstacle to the dependability position seems to be that the usual assumptions about failure modes and rates are inappropriate for survivable systems. It is far from clear that we come up with tractable models of survivable systems that include reasonably faithful failure modes and hostile capabilities.

Can we develop models that are reasonably independent of the details of failures modes, probabilities and correlations (given that in the context of survivable systems these are largely unknowable)? Unfortunately, given current engineering practice, such models are unlikely to be faithful to reality.

A system that is amenable to analysis in this sense will also probably be robust and resilient. That is, the characteristics that make a system easy to analyze are probably the same as those that make it survivable. Small, localized damage should result in small, localized degradation of service.

One possible response to the above observations is to try to develop guidelines for the development of robust, resilient systems and architectures, perhaps along the lines of the Abadi and Needham [1]. Whether such guidelines would ever be followed is another problem.

The need to reason about open, unbounded systems

Typically the systems we need to consider will be open and unbounded. Consequently we are attempting to establish properties of highly uncertain systems existing in an uncertain, shifting environment. Indeed the boundary between the system and its environment becomes blurred.

A possible reaction is take an agent-centric point of view and abandon attempts to prove global properties. Rather, try to show that from the point of view of individual agents the system provides certain standards of service. Thus the properties of interest should be derivable from certain assumptions about parts of the system under the agent's control, or at least trusted, along with minimal assumptions about the rest of the universe. This is akin to the approach taken in the analysis of security protocols in particular the strand spaces approach [6, 9].

Data-independence, abstraction and induction techniques have been used successfully to prove (conventional) properties of unbounded systems, [10]. It seems likely that similar techniques can be used for survivability properties.

The need to handle partial and faulty information about the global state

A survivable system will be made up of components with (local) interactions with the other components and the environment. Each component will need to take actions based on partial and probably faulty knowledge of the state of other components as well as the state of the (possibly hostile) environment. Note further, that hostile agents will probably be trying to subvert the communications mechanisms between the system components.

The need to reason about adaptive, self-stabilizing and emergent properties

It would appear to be necessary for any survivable system to incorporate emergent algorithms and mechanisms. Unfortunately we are still not very good at reasoning about emergent properties, in particular, our usual inductive, modular styles or reasoning are typically inappropriate.

Simulation tools like EASEL may help us develop better intuitions about emergent behaviors and help us develop techniques for reasoning about them [4].

The fact that survivability has so many aspects

The design and evaluation of a survivable system calls for consideration of dependability, reliability, security, adaptability etc. as well as efficiency, cost-effectiveness. It is thus hardly surprising that this has proved to be a challenging enterprise, especially given that it is far from clear that any clean separation of concerns is possible.

All those nasty economic, legal and policy obstacles

We will not dwell on these issues; they have been discussed at length by many people far better qualified than us. However, for completeness we say a few words.

Ordinary market forces tend to drive solutions towards a minimalist approach: just-in-time etc. As a result systems tend to be meta-stable: slight perturbations and they collapse. Achieving robust solutions tend to call for over-engineered solutions, but these are inevitably more expensive.

Building a survivable infrastructure is rather like working towards a cleaner environment: it requires a degree of responsibility and even altruism on the part of the people managing the components, qualities that can be rather thin on the ground in context driven purely by market forces. Combating this may require policy intervention and regulation to encourage the development and deployment of more robust survivable systems.

Another source of major problems is the universal insistence on using COTS components and the accompanying dogma that COTS components lead to cheaper and more effective solutions. Since we can have only uncertain knowledge of the behavior of COTS components and their interactions this leads to the issue mentioned above.

Is survivability survivable?

In the light of all this, is the pursuit of survivability the 21st century equivalent of the pursuit for the philosopher's stone? Perhaps survivability is beyond the reach of mere mortals and the best that we can hope for is longer mean times to fatality.

Can it survive attacks, failures and accidents? In other words, can survivability as a discipline survive these obstructions?

Is survivability itself an emergent computation? Is the approach for defining and understanding survivability the same that we seek for the systems we describe as survivable?

Conclusions

It is clear that the design and evaluation of survivable systems is an immensely challenging task. Even the question of establishing precise criteria for evaluating the survivability of systems is difficult.

Maybe the rather disappointing level of progress that we have witnessed is in fact roughly what we should expect from a new discipline. It might be instructive to look at the history of various other emerging disciplines such as security and artificial intelligence and compare their patterns of evolution. In security it took a couple of decades to reach the stage of having well-defined criteria for the evaluation of secure systems in the form of the "Orange Book" and even then these were controversial and have since been superseded [3]. Even now there is no consensus as to the correct formal definition of secrecy.

On the other hand, maybe survivability is more akin to alchemy.

We suggest that the way forward is to put aside the philosophizing and get our teeth into some realistic examples. Currently we are looking at models of "worm wars": worm systems seeking to propagate and attack a network along with "benign worms": seeking to defend the network. Both the worm system and the network are regarded as survivable systems with competing missions. This reflects rather nicely the two sided, game-theoretic nature of the problem. This work will appear in a separate paper [2].

Acknowledgement

The authors would like thank, amongst others, John McHugh, Tim Shimeall and Dave Fisher for fruitful discussions.

References

- [1] Abadi, M., Needham, R.: "Prudent engineering practice for cryptographic protocols"; IEEE Transactions on Software Engineering (1996), 22(1), pp 6-15.
- [2] Dietrich S. et al, "Worm Wars", to appear.
- [3] DOD 5200.28-STD, Department of Defense Trusted Computer System Evaluation Criteria (TCSEC) ("The Orange Book"), National Computer Security Center, December 1985.
- [4] Fisher D., "Design and Implementation of Easel, a language for simulating highly distributed systems", at <http://www.cert.org/research/>.
- [5] Jacob, J. "Basic theorems about security". Journal of Computer Security, 1(4), pp 385-411, 1992.
- [6] F. Javier Thayer Fabrega et al. "Strand spaces: Why is a security protocol correct?", Proceedings of the 1998 IEEE Symposium on Security and Privacy, pp 160-171. IEEE Computer Society Press, May 1998.
- [7] Ryan P. Y. A., "Mathematical Models of Computer Security", Proceedings of the FOSAD 2000 Summer School, to appear Springer LNCS 2172.
- [8] Ryan P. Y. A. et al, "Survivability, the Universe and Everything", to appear.
- [9] Ryan P. Y. A. et al, "Modelling and Analysis of Security Protocols", Pearson 2001.
- [10] Broadfoot P et al, "Automating data-independence", ESORICS, LNCS 1895, Springer 2000.