Consilience of Safety, Security and Resilience

Emil Lupu, Luca M Castiglione

Consilience of Safety, Security and Resilience

Emil Lupu and Luca M Castiglione
Imperial College London, UK
{e.c.lupu, l.castiglione}@imperial.ac.uk

I. Introduction

Modern safety-critical systems must be safe, secure, and resilient. However, the intersections of these properties are the source of contention and sometimes misunderstandings. Academic publications (e.g., [5]–[7]), and standards (e.g., [3]) discuss the *conflicts* arising between safety and security requirements. Broadly, this happens because security restricts behaviours that can occur, whereas safety relies on certain behaviours occurring. So, the two will be in tension when different requirements (e.g. confidentiality and availability) are considered separately. Similarly, most researchers agree that security is needed for safety, yet at the same time, that security is not fully achievable due to system size and complexity. *Resilience* is usually considered as minimising the loss of function over time. When faced with an adversary, this can be ensured by: a) improving security b) increasing redundancy and/or capacity c) enabling faster recovery, or d) a combination of these. Choosing the optimal strategy, is the fundamental challenge in resilience research. Yet the adaptation required by resilience is in tension with the rather static and lengthy safety assurance processes.

We present a framework for simultaneously reasoning about safety, security and resilience, discuss the respective analyses and their intersections drawing lessons on the salient research gaps. Our framework aligns with consequence driven cyber-informed engineering (CCE) [4] and we favour an informal descriptive approach over one based on definitions, ontologies, or epistemology.

II. Characteristics and Trends

Regardless of how one defines a *system* and whether the application context refers to (critical) infrastructures or autonomous systems, some common characteristics and trends emerge. Most systems are increasingly *cyber-physical*, i.e., they interact with physical phenomena and are affected by the physical environment, which can itself be adversarially subverted. Security must thus consider both the "cyber" and the "physical" aspects of the system. Most systems are increasingly *complex* and comprise multiple interacting components, so unexpected interactions can occur. This complexity makes analysis difficult and gives rise to *cascading effects* whereby changes to the input or behaviour of a component propagate to its dependent components. Finally, the number of *system components* is increasing, leading to broader attack surfaces. Interconnections between components enable attacks to propagate, whilst component interdependencies lead to cascading effects that propagate the effects of (malicious) actions, i.e., the *impact* of the attack. In this context, it is unrealistic to guarantee "security" and compromised components must be assumed to behave in any way the attacker chooses.

We are grateful for financial support and helpful advice to the Defence Science and Technology Laboratory (DSTL) – (EPSRC EP/X037452/1).

III. A framework for system modelling

We have found it useful to adopt a layered approach and distinguish between the *network topology* layer representing the interconnections between the devices on which the components are hosted and the *functional* layer representing the interdependencies between the components (Figure 1). The topology allows reasoning about how an attacker can *compromise* the system by moving laterally and acquire more privileges, i.e., a *position* in the system [1]. In contrast, the functional layer enables reasoning about the *impact* of *attack actions* (which perturb components' operation) and their cascading effects. To make it easier to reason about system compromise, we introduce an intermediate layer representing the *attack graph*, i.e., the possible compromise pathways and privileges acquired by the attacker. The *position* of an attacker is the set of nodes visited by an attack along a path. The set of all paths represents the positions the attacker can achieve.

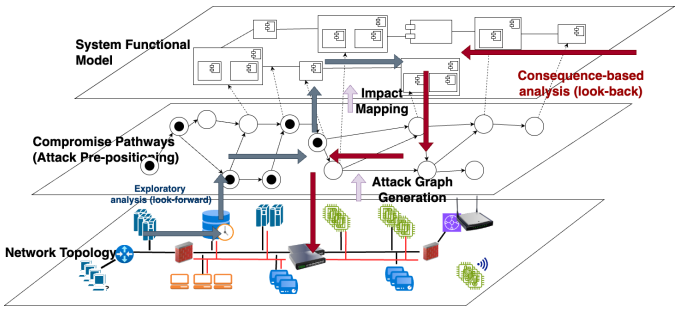

Fig. 1. A framework for safety, security and resilience analysis

Safety and (more generally) Consequence-based Analysis. Consequence-based analysis (including safety analysis) is conducted in the functional layer, starting with the consequences of interest and considering what might cause them. In safety, the consequences are typically losses, which may also include economic losses, in addition to injury or threats to life. The analysis recurses through the component interdependencies in the functional layer in a *"look-back"* fashion, i.e., identifying

the inputs that cause a particular output. When considering security, this requires identifying the attack actions leading to the effects, the positions that confer the privileges necessary to perform those actions and the attack paths leading to them. A consequence-based approach has the advantage of significantly reducing the range of inputs that must be considered. From a security perspective, only the positions that can lead to losses (and the attack paths leading to them) need to be analysed. However, identifying *all* the possible causes, i.e., *all* the combinations of inputs that lead to an output, across *all* cascading dependencies is difficult due to non-linearities, interaction with physical phenomena, feedback, etc. When the causes are *failures*, we have experience of which values to consider. When the input is maliciously crafted, the entire domain of values must be considered.

Security Analysis. Security analysis typically comprises different aspects. *Threat modelling* in the functional layer helps identify malicious (attack) actions that an adversary may perform to affect the system. This abstracts out how the attacker may acquire the *position* needed to perform those actions. The latter analysis is conducted through different techniques, including vulnerability analysis, attack-graph analysis, and pen-testing. If consequence-based analysis *looks back* from the consequences, here the analysis *explores*, i.e., *looks forward* from a point of access considering the consequences of attack steps. This is somewhat easier, but results in numerous potential attack-paths, too many to manually enumerate and analyse. Where consequence-based analysis meets the exploratory security analysis is rarely considered and usually not formalised, leading to misguided outcomes when the presence of weakenesses and vulnerabilities is directly linked to the effects of attack actions. Failing to differentiate between the compromise position and performing an attack action is an important omission. Often a threat actor seeks to acquire and persist a position in the system, without immediately impacting system operation, e.g. [2]. Informally, the attacker's appetite to obtain a capability to observe and cause damage is different from their appetite for causing the damage itself. Any risk assessment, whether a probabilistic or a time-based calculation, needs to consider the two aspects separately. Our proposed framework allows to make this distinction; we consider separately the attack paths (solid arrows) and the impact attack actions (dotted arrows) in Figure 1. Our prior work, [1] combines security and safety analysis in a novel way to systematically identify attacks which, in conjunction with cascading effects, lead to safety violations.

Resilience Analysis aims to determine and minimise the loss of function over time when perturbations occur including, failures, performance degradations, or malicious attacks. Our framework shows why this is complex when considering cyber-attacks. At each time point, an attacker having acquired a foothold in the system, can choose whether to continue compromising the system (remaining in the attack graph layer) or trigger an impact-generating attack action. This leads to a combinatorial explosion of *attack scenarios*. Our prior work shows that these alternatives have a significant effect on the resilience strategy and reveal hidden factors that influence strategic choices [8], [9]. For example, the *dwell time* of the attacker between the moment of initial compromise and the moment of detection alters the relative effectiveness of compromise isolation and system recovery. For non-adversarial threats, resilience analysis considers perturbations such as failures and damage from extreme weather, which can be modelled according to past history. For adversarial attacks, fewer assumptions can be made and comprehensive simulations and emulations of attack scenarios are needed.

IV. Conclusions

Systems must be simultaneously safe, robust and resilient. However, the consilience of these properties requires the right framework. We have outlined a proposed framework that separates the compromise part of the attack (acquiring a position) from the attack actions and their cascading impact.

We have shown that it is useful to abstract network reachability and attack step chaining as attack graphs – an established concept in security but rarely used in safety or resilience. We have further shown that safety and more generally, consequence-driven analysis, "looks back" whilst security and resilience "looks forward". Across layers, this corresponds to a top-down vs. bottom-up analysis.

Techniques for bi-directional chaining to combine these analyses are to our knowledge unexplored. Similarly, search strategies and heuristics in the exploration of the design space that efficiently lead to solutions simultaneously satisfying safety, security, and resilience remain to be developed. However, there is convergence between the different communities and the manual reconciliation of security, safety and resilience is so labour intensive that there is a significant demand for new automation techniques.

References

[1] L. M. Castiglione and E. C. Lupu. Which attacks lead to hazards? combining safety and security analysis for cyber-physical systems. *IEEE Trans. Depend. Sec. Comput.*, 21(4):2526–2540, 2023.

[2] CISA. State-sponsored actors compromise and maintain persistent access to us critical infrastructure, 2024.

[3] EUROCAE. Ed-203a - airworthiness security methods and considerations, 2018.

[4] S. Freeman, C. St Michel, R. Smith, and M. Assante. Consequence-driven cyber-informed engineering (cce). Technical report, Idaho National Laboratory (INL), 2016.

[5] S. Kriaa, L. Pietre-Cambacedes, M. Bouissou, and Y. Halgand. A survey of approaches combining safety and security for industrial control systems. *Reliability Engineering & System Safety*, 139:156–178, 2015.

[6] S. M. Nicoletti, M. Peppelman, C. Kolb, and M. Stoelinga. Model-based joint analysis of safety and security: Survey and identification of gaps. *Computer Science Review*, 50:100597, 2023.

[7] C. Ponsard, J. Grandclaudon, and P. Massonet. A goal-driven approach for the joint deployment of safety and security standards for operators of essential services. *J. of Software: Evolution and Process*, page e2338, 2021.

[8] J. Soikkeli, G. Casale, L. Muñoz-González, and E. Lupu. Redundancy planning for cost efficient resilience to cyber attacks. *IEEE Trans. Depend. Sec. Comput.*, 20(2):1154–1168, 2022.

[9] J. Soikkeli, L. Muñoz-González, and E. Lupu. Efficient attack countermeasure selection accounting for recovery and action costs. In *Conf. on availability, reliability and security (ARES)*, pages 1–10, 2019.