Assumptions in computer science: mere mathematical hypotheses, or representations of the physical world?

Paulo Esteves-Veríssimo

paulo.verissimo@kaust.edu.sa

https://cemse.kaust.edu.sa/people/person/paulo-verissimo

https://rc3.kaust.edu.sa

KAUST, CEMSE, RC3 (Resilient Computing and Cybersecurity Center)

The praxis of algorithm or mechanism design, especially in concurrent or distributed systems, mandates clear statements of the assumptions underlying the design, such as topology, dimension, synchrony, performance, threats, etc. The design --- say, of an algorithm whose behaviour is defined by a set of properties (safety, liveness) --- is then shown correct by demonstrating it does secure those properties, given the assumptions. From a mathematical viewpoint, we are done. We never ask: are the hypotheses valid?

If we wish the algorithm to have any real world impact, it might be advisable to define an 'abstract system' credibly materializing the assumptions made, where the protocol implemented from the algorithm will run correctly. However, this apparent detail completely changes the perspective: in essence, the protocol correctness becomes conditional to the likelihood of the assumptions being met, in such a system.

Suddenly, we must look at our assumptions from a physics viewpoint: How do I achieve perfect failure detection in an asynchronous environment? Why would an attacker compromise certain units and not others in an otherwise arbitrary failure environment? We should not be mistaken about these being "implementation details". They can and should be addressed as 'systems theory' problems, related to the 'substance' and the 'robustness' of assumptions --- i.e., the coverage of the mapping of the 'abstract system' onto the physical world.

The rise of malicious threats to systems has been showing the importance of this argumentation. The practice of accepted deviations of some assumptions from physical reality, in face of accidental (stochastic) threats, completely crumbled in face of malicious (intentional) threats, exposing their lack of substance and/or robustness. Paraphrasing and extending my colleague and friend Fred Schneider's quote some years ago: «Every [non-substantiated] assumption is a vulnerability».

In the talk, I will delve into manifestations of the problems above, and approaches to solve them in a satisfactory way. Closing these gaps implies effort on some angles, and I will single-out two: (i) system architecture and design; (ii) modeling and verification. I will be discussing: system awareness (topology-, context-, hybridisation-), logical vs. physical centralisation, trust vs. trustworthiness); system-level impossibility results, lower bounds and safety predicates (coverage-stability, no-contamination, exhaustion-safety) denying substance to some commonly made assumptions --- or providing guidance to achieve it.