

# Security

## Coming of Age

*Stressing the importance of threat models.*

**A**S FOR MANY colleagues of a similar age to mine, my path into computer security research was all but academic. In the late 1990s, there were very few academic research centers on security, and even less dedicated courses or degrees, especially outside the U.S. We—many of us, at least—came from a hacking background, experimenting with things such as vulnerabilities and exploitation on our own. We would find our brethren in obscure alleyways of the Internet, and then browse through e-zines and (if lucky) attend hacker meetups with a score of attendees. I remember fondly the feeling of attending DefCon for the first time, 20 years ago, and seeing a few thousand kindred souls together.

It should not come as a surprise that for the “hackademics,” as a colleague once half-jokingly defined us, offensive security research has a definite thrill. In a discipline that lacks a fundamental, unified theory of how to build “secure things,”<sup>1</sup> and where in fact most properties are defined in negative terms (how “not to build”), this makes rational sense. After all, we define robustness of encryption based on resilience to attacks: We routinely first propose attacks, and then offer mitigations. Even in the applied, corporate world, we use penetration testing and red teaming exercises to assess security level. The strictest security evaluation standards, such as the Common Criteria, define security on the basis of resilience to attack attempts. As the saying goes, in security defense



is the child of offense. Not the other way around.

But this is just partially true. The fact is, those with a similar background to mine would not want it any other way. We take pride and delight in clever hacks and bypasses. We en-

joy reading a clever exploitation write-up, in the same way a mathematician enjoys a brilliant demonstration: an intellectual pleasure more akin to appreciation of art than anything else no matter how narrow the applicability, no matter how obscure the trick (in fact, the more, the better), no matter the prerequisites or the real-world impact. Capture-the-flag (CtF) exercises and competitions are somehow exemplar in this: There is absolutely no comparison between the ones we used to play in 2003–2004 and those my students play today. Besides becoming a mainstream tool for education and a competitive form of entertainment (as opposed to a curiosity for a small group of nerdy students worldwide),

**As the saying goes,  
in security defense is  
the child of offense.**

acm

## Advertise with ACM!

Reach the innovators  
and thought leaders  
working at the  
cutting edge  
of computing  
and information  
technology through  
ACM's magazines,  
websites  
and newsletters.



Request a media kit  
with specifications  
and pricing:

**Ilia Rodriguez**  
+1 212-626-0686  
acmm mediasales@acm.org

acm

media

they have become amazingly complex and challenging puzzles and brain teasers, requiring a level of skill and understanding that amazes and thrills me. At the same time, they have moved incredibly far from (boring) real-world security vulnerabilities or penetration-testing exercises, as more than one disappointed former student has confessed to me after joining a real-world corporate red team.

What's more, hackademics love the practical demonstration. We live to impress with the stunt: making the ATM spit bills on the floor, live on stage; shooting a video of the robot going rogue; driving a car off the road in front of the press. Subsequently, our conferences (not just the academic ones, but also the industry ones, such as Black Hat, and the community ones, such as DefCon or BSides) reward attack research more than defense.

In the meantime, while our community's interests somehow grew more and more esoteric, the world around us changed. Even the words I used in this nostalgic incipit are not really used anymore or have been thwarted. "Hacker" definitely means something different for the general public, and our discipline has grown both in numbers and in breadth with the transition to the broader concept of "cybersecurity," mirroring the growing, ubiquitous importance of computer systems in our hyperconnected world. Computers chart safe paths for aircraft through the skies; manage the distribution of electricity, the lifeblood of modern society; and help us throughout our day, both evidently and behind the curtain. And in this growth and transformation of both cybersecurity and the broader computing field, what was once so obscure that it lacked recognition even in academia is now the subject of front-page news.

People now perceive cybersecurity as an issue of growing importance—in particular when linked to cyber-physical systems that can affect their lives and safety;<sup>5</sup> thus, they look for information, understanding, and appropriate guidance about it. And—broadly speaking—we should be the ones providing it, but our discipline, research, and more importantly,

our method to communicate this research is set up for anything but.

Most cyber-physical security research, for instance, is aimed at the exploitation of specific vulnerabilities, and at stunt hacking. See? I can steer the car via remote control! Watch! I can make the plane fly sideways. It is, once more, the fascination of the hack. It is what we love to do and see, but it is painfully far from what is needed to teach the public to think sensibly about security issues.

If I had to choose one lesson I learned and verified, again and again, over the past 25 years, it is that at the heart of every real major security issue, as well as every overhyped or misguided vulnerability announcement, was a common root: a failure in threat modeling.<sup>4</sup> The only way to reason sensibly about the security of a system, the closest thing we have to a method in the madness of this field, is to start from a realistic and complete model of the threats and the attack surface of the system. This may seem a very basic, even naive, observation: After all, threat models are one of the few literal "textbook materials" we can teach in a university course on security. However, careful observation of past and current research and events shows a painful lack of methodical application of this basic conceptual tool.

While reading the vast amount of literature on the (in)security of vehicles, for instance, it would be extremely easy to fall prey to security nihilism and desperation: everything that could be broken is apparently

**Most cyber-physical security research is aimed at the exploitation of specific vulnerabilities, and at stunt hacking.**

broken, and most things are extremely challenging or impossible to fix in the near term. On the other hand, cyberattacks in the automotive domain are anecdotally few (if we exclude those aimed at stealing vehicles). A sensible analysis of the threat model explains that while vulnerabilities are extensive (and should be addressed in the long term), the potential attacker goals are quite specific: influence or disrupt the safe operation of the vehicle, extract PII from it, or steal it. Of these, the only one supported by the existence of threat actors against the general population is the last one (which is the prevalent one). Safety-threatening attacks will be an arising issue when fleets of autonomous vehicles will roam around, but they are not really a viable scenario right now. By assessing the threat model sensibly, and by communicating our research accordingly, we can inform the public and drive a rational, risk-driven approach, as opposed to knee-jerk reactions and panic,<sup>3</sup> and ensure future technology, such as autonomous driving, will be safe and secure when ready for deployment in a few more years.

Of course, this approach requires us to be more humble and more precise in specifying and communicating the limitations of our research; and in turn, it requires our community to restructure expectations and reward models accordingly in our publications and conferences. An example in my mind is aviation. Some researchers focused on the security weaknesses of communication protocols used in modern aviation, most notably ADS-B.<sup>2</sup> And they properly outlined authenticity and availability issues, stemming from a design decision of not adopting any encryption or signing for the protocol messages. However, the (potential) impact of such vulnerabilities is limited, if we consider how the data is used, the layers of (non-digital) operational safeguards in how planes are flown, and in general if we once again appropriately build a threat model and consider everything that would be needed to carry out a successful attack. Does this detract from the importance of the research, or does it rather help drive it in appropriate directions, focusing


## Threat models are important, because they abstract from the technical details of a specific hack.

it on the many relevant issues in the aviation domain? We can definitely study and improve the security of avionics without overhyping the results, or implying non-existing immediate threats to the safety of passengers and crews. Security vulnerabilities can be interesting and worthy of attention without being necessarily world-shattering.

Appropriate threat modeling of complex systems, and particularly of cyber-physical systems of societal or strategic significance, also requires true interdisciplinarity. Threat models of transportation systems, for instance, require the input of safety and transportation engineers, but also of terrorism experts and political scientists, not to mention environmental protection experts. Modeling appropriately the vulnerabilities of an aviation component or protocol may require the expertise of pilots and air traffic controllers. Way too often cybersecurity experts oversimplify the external drivers of risk, confusing the ends with the means. The specific detail of how can a SCADA controller be compromised pales in comparison to the systemic threat posed to the society by a catastrophic failure in the electrical grid, and in turn if such a disaster can be avoided through appropriate system design, the potential compromise of a controller does not really matter as much as we would like to think. In fact, thinking as hackers, and not as engineers, can bring us to identify ignored weaknesses and vulnerabilities, but it is only by thinking systemically that we can help make systems resilient, and society safer. In many ways, we should aim to make cyberattacks

(and ourselves) much less relevant.

Threat models are important, because they abstract from the technical details of a specific hack: They tend to be reasonably understandable even by non-specialists, and definitely within reach of the other practitioners and engineers we desperately need to interact with. They allow us to place our shenanigans in context and to see the bigger picture of what is actually relevant. Appropriate threat modeling is key to understand and fix the underlying problems, instead of being distracted by the cyber-emergency of the moment, be it denial-of-service campaigns related to political events or ransomware groups running amok. Last but not least, threat models allow us to communicate appropriately our discipline to our colleagues, and to fit it into the bigger picture of systems engineering.

Threat models sound pretty boring—but tend to be absolutely fundamental, like most things related with adulthood. Cybersecurity has come of age, and as much as we can (and *should!*) keep being delighted by the beauty of the hack, it is time we started behaving and communicating as expected by professionals in a society-relevant field. 

### References

1. Bratus, S. et al. Why offensive security needs engineering textbooks: Or, how to avoid a replay of "crypto wars" in security research. *login Usenix Magazine* 39, 4 (2014).
2. Costin, A. and Francillon, A. Ghost in the air (Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices. *Black Hat USA* (2012).
3. Longari, S. et al. A secure-by-design framework for automotive on-board network risk analysis. *2019 IEEE Vehicular Networking Conference (VNC)* (2019), 1–8; DOI: 10.1109/VNC48660.2019.9062783
4. Shostack, A. *Threat Modeling: Designing for Security*. (First edition). Wiley Publishing, 2014.
5. Zanero, S. When cyber got real: Challenges in securing cyber-physical systems. *2018 IEEE Sensors (2018)*, 1–4; DOI: 10.1109/ICSENS.2018.8589798

**Stefano Zanero** (stefano.zanero@polimi.it) is a professor at NECST Laboratory, Politecnico di Milano, Milan, Italy.