

# Cybersecurity Education and Formal Methods

James H. Davenport<sup>1</sup> and Tom Crick<sup>2</sup>

<sup>1</sup> Department of Computer Science, University of Bath, UK  
jhd@cs.bath.ac.uk

<sup>2</sup> Department of Computer Science, Swansea University, UK  
thomas.crick@swansea.ac.uk

**Abstract.** Formal methods have been largely thought of in the context of safety-critical systems, where they have achieved major acceptance. Tens of millions of people trust their lives every day to such systems, based on formal proofs rather than “we haven’t found a bug” (yet!); but why is “we haven’t found a bug” an acceptable basis for systems trusted with hundreds of millions of people’s personal data?

This paper looks at some of these issues in cybersecurity, and the extent to which formal methods, ranging from “fully verified” to better tool support, could help. More importantly, recent policy reports and curricula initiatives appear to recommend formal methods in the limited context of “safety critical applications”; we suggest this is too limited in scope and ambition. Not only are formal methods needed in cybersecurity, the repeated and very public weaknesses of the cybersecurity industry provide a powerful motivation for formal methods.

**Keywords:** Formal methods, cybersecurity, curricula

## 1 Introduction

Formal methods, when they have been thought of at all, have been largely thought of in the context of safety-critical systems, where they have achieved major acceptance in what is, alas, an unsung area of software development. Tens of millions of people trust their lives every day to such systems, but nearly all are unaware of these systems, and the extent to which they are enormous successes. Even people “who ought to know better” don’t. One of the authors quoted the Paris Métro Ligne 14 performance figures (software shipped in 1999 and no bugs reported [1]) to a major figure in the commercial software industry, to be told that he was lying, as this was utterly impossible.

Formal methods *ought* to be much more widely used in the cybersecurity industry. This is much more visible (because it has many conspicuous failures) than the largely invisible safety-critical industry. However, formal methods are not currently widely adopted here, and hence there is tremendous scope for growth and adoption of formal methods. In addition, as the cybersecurity industry and its failures are much more visible, emphasising the relevance of formal methods to the cybersecurity industry should encourage more interest at universities.

## 2 Cybersecurity

Cybersecurity<sup>3</sup> failures abound, and the number of people that can be affected by even a single failure is amazing — 148 million for Equifax [2] and probably more for the Starwood<sup>4</sup> breach: a number [3] “downgrades” to 383 million. The financial costs can be substantial: bankruptcy in the case of American Medical Collection Agency [4] and a provisional £183M fine for British Airways [5]. These problems have attracted attention at the highest scientific levels [6].

There are many reasons for cybersecurity failures, and even a given failure may have multiple causes. For example, the U.S. Government investigation [7] into Equifax states “Equifax’s investigation of the breach identified four major factors<sup>5</sup> including identification, detection, segmenting of access to databases, and data governance that allowed the attacker ...”. However, none of these would have been triggered had it not been for the original bug in the Apache code [9], which was of the well-known (Number 1 Application Security Risk in [10]) family of “Injection” (or “Remote Code Execution”) attacks, and which would probably have been detected by an automatic taint analysis tool such as [11].

Though attributing causes at scale is difficult, a well-known textbook [12] claims that about 50% of security breaches are caused by coding errors. Hence it behoves security practitioners to look seriously at coding errors, while recognising that this is only one facet of the problem. This is taken up by the Payments Card Industry in [13], essentially the only world-wide mandatory security standard, in two sub-requirements of “Requirement 6: Develop and maintain secure systems and applications”.

**6.5** Address common coding vulnerabilities in software-development processes as follows:

- Train developers at least annually in up-to-date secure coding techniques, including how to avoid common coding vulnerabilities;
- Develop applications based on secure coding guidelines.

**6.6** For public-facing web applications, address new threats and vulnerabilities on an ongoing basis and ensure these applications are protected against known attacks by either of the following methods:

- Reviewing public-facing web applications via manual or automated application vulnerability security assessment tools or methods, at least annually and after any changes;

---

<sup>3</sup> The precise definition of cybersecurity is debatable: we can take it as failures of security, generally defined as “preserving the CIA — Confidentiality, Integrity and Availability” of digital information, where computer system played a critical part in the failure.

<sup>4</sup> Generally called “Marriott”, but in fact due to the Starwood chain before Marriott took it over.

<sup>5</sup> In military parlance, Equifax is being found not to have “defence in depth”. Defence in depth is certainly valuable: [8] described how Google was saved from the consequences of an ‘awesome’ attack on gmail by defence in depth. But the front line is still the first defence: in this case correct code.

- Installing an automated technical solution that detects and prevents web-based attacks (for example, a web-application firewall [WAF]) in front of public-facing web applications, to continually check all traffic.
- \* The authors note that many sites go for the second option, as it's easy to “tick the box” Indeed [14] finds that 34% of customers regard compliance as the main function of their firewall, rather than security. Apart from the fact that firewalls can be misconfigured, as in the Capital One attack [15], they can also be placed in “detect-only” mode and ignored<sup>6</sup>.
- A further challenge to the PCI DSS model is provided by the modern “Magecart” attacks, such as that against British Airways [17]: here the JavaScript that is downloaded to the customer’s browser is corrupted, and this leaks the data directly, without going near the WAF. This raises the challenge of JavaScript verification §5.4.

It is noteworthy that, despite apparently insisting on secure coding in 6.5, they require the additional defences in 6.6, realising that *errare humanum est*, and the 6.5-developed code may not actually be secure. Is it possible (the authors think so, but the experiment has yet to be performed) that adding formal methods to 6.5 would render 6.6 redundant, or at least mean that 6.6 should be restricted to finding design errors, rather than debugging 6.5 failures? Full formal verification of a complete system should certainly suffice.

Complete formal verification is the only known way to guarantee that a system is free of programming errors. [18, describing seL4: a verified operating system]

Such a verified operating system has been used in medical devices, but probably not sufficiently widely, as 500,000 already-fitted pacemakers have had to be upgraded through security weaknesses [19], and insulin pumps are also vulnerable [20], as many others [21, 22]. See [23] for a recent update on seL4. However, most of us do not have the opportunity to start from scratch, and have to live on top of imperfect, unverified systems, interoperating with other systems via large, generally unverified, or at least under-verified, protocols, such as TLS [24].

### 3 Agile versus Secure

“Agile Development” [25] is a major theme in software development. Mark Zuckerberg can be said to have taken this theme to the extreme in 2009.

“Move fast and break things” is Mark’s prime directive to his developers and team. “Unless you are breaking stuff,” he says, “you are not moving fast enough.” [26]

---

<sup>6</sup> [16]: “Midsized and small companies frequently install WAFs just to satisfy a compliance requirement. They don’t really care about practical security, and obviously won’t care about maintaining their WAF. ” This is backed up by [14], whose survey states “ 43% run their WAF in alert-only mode!”.

In both safety-critical and security-conscious programming, “breaking things” comes with a very high price. Aeroplanes can’t be uncrashed, and data can’t be leaked.

The problems with using “Agile” methods in security are well-documented, at practitioner level, e.g. a recent “Security + Agile = FAIL” presentation [27], in many theoretical analyses as well as the interview-based research in [28] for small teams and [29] for large multi-team projects. Both mention team expertise in security as a significant problem.

**From [28]** The overall security in a project depends on the security expertise of the individuals, either on the customer or developer side. This corresponds to the agile value of “individuals and interaction over processes and tools” [25, Value 1].

**From [29]** The interviewees generally agree that more could be done to provide security education and training to employees. Without prompting, several interviewees mentioned training as an important factor for increasing security awareness and expertise.

It is very hard to take security seriously in this setting.

**From [28]** security “is only of interest [to the customer] when money-aspects are concerned”.

**From [29]** One Test Manager articulated his team view that “security is not currently seen as part of working software, it only costs extra time and it doesn’t provide functionality”. With less focus on providing extensive (security) documentation typical for agile, ineffective knowledge sharing between security officers and agile team members is especially problematic.

**From [30]** (A more general survey, but many papers surveyed were “Agile”) “Security is often referred to as a NFR [non-functional requirement] in that it is expected to be included as part of high quality code development, but is rarely listed as an explicit requirement. As a result, developers prioritise security below more-visible functional requirements or even easy-to-measure activities such as closing bug tracking tickets.”

It would be tempting to conclude that “Agile” and “Secure” are, or at least are close to being, mutually contradictory. But there has been some analysis of the same apparent contradiction in the safety-critical industry [31]. Other than “Embedded Systems”<sup>7</sup> [31, §3.6], this analysis of the problems is fairly close to the practitioner view in [27], and we could reasonably ask what lessons could be carried across.

## 4 The Need for Tools

There are two key points.

---

<sup>7</sup> Actually, Embedded Systems are a comparatively neglected, but important, cyber-security area. See, for example, [32] for a description of a pervasive design fault in the “home security” market.

**From [31, §4.1]** Strong static verification tools tend to complement (not replace) human-driven review<sup>8</sup>. The tools are very good at some problems (e.g. global data flow analysis, theorem proving) where humans are hopeless, and vice versa. If we do the static verification first, then we can adjust manual review processes and check-lists to take advantage of this.

**From [31, §6]** The sixty-four-million-dollar-question, it seems, is how much “up-front” work is “just right” for a particular project. We doubt there’s a one-size-fits-all approach, but surely the answer should be informed by disciplined requirements engineering of non-functional properties (e.g. safety, security and others) that can inform the design of a suitable architecture and its accompanying satisfaction argument.

Facebook grew, security (and “product quality” in general: it is not clear whether security was the main driver here) became more important, and by 2014 Zuckerberg had changed his views.

“Move fast with stable infrastructure.” It “may not be quite as catchy as ‘move fast and break things,’” Zuckerberg said with a smirk. “But it’s how we operate now.” [34]

One might think his views were converging with the views of [31]. However, the Heartbleed story [35] should remind us that the fact that a modification “has no new security considerations” *as designed* [36] doesn’t mean that an implementation of that idea has no new security considerations. Hence the call in [31, §4.1] for strong static verification tools. Such tools are generally seen as expensive and slowing down the development process, but [37] shows that they need not be. In particular, they show that, for a real application (890,000 physical lines of Ada code), the cost of incremental verification can be reduced from “nightly” to “coffee”, and hence can reasonably form part of a continuous integration toolchain, as is done at the company studied in [37]. Readers might comment that their own applications are not in Ada, but [38, §5.6] discusses mixed-language programming, especially with C. A similar point is made in [39], describing the Infer tool running on Java/Objective C/C++, where moving from overnight reporting to near real-time reporting moved the fix rate from 0% to 70%.

That these techniques are reaching the mainstream of cybersecurity can be seen from Amazon Web Services adoption of them [40], Google [41], Facebook [39], and the recent DefectDojo release by OWASP [42].

## 5 The Scope of Tools and Formal Methods

There is a substantial range of tools, and degrees of formality, and [31, §6] is probably correct in saying “We doubt there’s a one-size-fits-all approach”. At one extreme, there are the humble, but still surprisingly effective, `lint` and its equivalents, looking, essentially, for dangerous or dubious, though legal, syntax.

---

<sup>8</sup> A point made in the context of XP and Agile in 2004 [33].

## 5.1 Ada and SPARK

At the other extreme, there are languages, such as the SPARK Ada subset [38] designed with verification in mind and heavily employed in the safety-critical sector such as railways and air traffic control, which can also be deployed for demanding secure applications, such as an RFC4108-compliant secure download system for embedded systems [43].

## 5.2 C/C++

There is, however, a large middle ground between these two extremes. Even if the application is required to be in C or C++, there is a lot to be said for sticking to a safer (even if not provably safe) subset of the language *and associated libraries*, such as eschewing `strcpy` in favour of `strncpy`. This can often be enforced by static verification tools. We note that Google’s “Zero Day” project reports [44] that 68% of all such zero-day exploits (i.e. exploits discovered in the wild first) were caused by memory corruption errors, and Microsoft report a very similar story [45].

There is a good survey of such subsets and standards in [46, Appendix F]. As that notes, the ISO standard for secure C coding [47] has the unusual (for this middle ground) but important concept of “taint analysis” (as in [11]): input data should be considered “tainted” until it has been sanitised. This is particularly important for network-oriented applications, where it is natural for the programmer to believe that the other party is behaving correctly (as in **Heartbleed** [35]).

After this paper was presented, [48] appeared. That paper’s authors had formally proved properties of non-trivial parts of Amazon’s core C library. “We proved that key components of AWS C Common are memory safe, i.e. do not suffer from issues such as buffer overflow, use after free, or invalid pointer dereferences. Memory safety errors are routinely listed among the most critical security concerns by industry groups monitoring CVEs”.

## 5.3 Java

Closer to the SPARK Ada end of the spectrum we find Safety-Critical Java [49]. The authors do not have enough experience with this to comment directly. However, the Java ecosystem (Stack Overflow etc.) is far from security-aware [50]. The fact that an application is in Java doesn’t mean it’s free from security coding errors: see [51] for a recent example.

There is a static analysis security tool for Java described in [11]. As with [47], this has “taint analysis” as its major feature, and at the time it spotted some significant-seeming problems.

## 5.4 JavaScript

JavaScript is a particular problem for Security. There are some verification tools, e.g. GATEKEEPER as described in [52]. However, even if it were possible to

guarantee a particular piece of stand-alone JavaScript, that is not how the current paradigm operates. As [53] writes:

Much of the power of modern Web comes from the ability of a Web page to combine content and JavaScript code from disparate servers on the same page. While the ability to create such mash-ups is attractive for both the user and the developer because of extra functionality, code inclusion effectively opens the hosting site up for attacks and poor programming practices within every JavaScript library or API it chooses to use.

Though not explicit in this statement, an additional weakness is that this combination is *dynamic*. The obvious solution would be some kind of sandboxing of the external resources relied upon, but the nature of JavaScript makes this difficult. [54] describe one such sandboxing, but it only works for a subset of JavaScript and relies on a combination of filtering, rewriting and wrapping to guarantee security. That it can do so at all is a remarkable feat of formal methods, given that previous attempts such as Facebook’s FBJs have subtle flaws [55], and that the formal semantics of JavaScript being relied upon are very much a piece of reverse engineering.

In fact the dynamic loading from multiple sites is often not good for performance, and web performance engineers recommend tools to bundle the pages: this could usefully be combined with the sort of protection described by [54].

An alternative solution is suggested by Google, who are introducing a form of taint analysis into Chrome [56] through run-time typing. When enabled, this means that the 60+ dangerous DOM API functions can only be called with arguments whose type is that emitted by `TrustedTypes` functions. Google expects that these functions would be manually verified, but this does open the door to formal verification of *certain* security policies in what is currently a very challenging environment for formal methods. However, these checks can be easily fudged, and the authors foresee examples of this on StackOverflow analogous to the `csrf().disable()` “suggestion” described below in point 3.

## 6 Education

[13, Requirement 6.5] called for education of developers. Education of mainstream programmers, as opposed to cybersecurity specialists, in cybersecurity has been neglected until recently, and this neglect has been lamented as far as the Harvard Business Review [57]. Developments in professional accreditation are changing this [58]. However, there are limitations, even beyond *errare humanum est*, in relying on education.

1. There is experimental evidence that both trained students [59] and professional developers [60] will ignore security considerations unless *explicitly* instructed to take them into account. Lest this be thought to be a purely academic exercise with little relevance to the real world, consider the recent ¥55M password problem described in [61].

2. There is field evidence that explicit requirements such as [13] are ignored in practice, e.g. the Forever 21 breach [62], or Macy’s [63]. They may also not be communicated down the software supply chain, as in the Ticketmaster case [64].
3. Many educational resources, both formal textbooks [65] and informal resources such as Stack Overflow [66], pay very little attention to security, and indeed can be positively harmful. The discussion in Stack Overflow (analysed in [50, §4.3.1]) of cross-site request forgery (CSRF — this was in the OWASP top 10 in 2013 [67], but dropped from [10] “as many frameworks include CSRF defenses”) is especially worrying. By default, Spring implicitly enables protection against this. But all the accepted answers to CSRF-related failures simply suggested disabling the check. There were no negative comments about this, and indeed a typical response is “Adding `csrf().disable()` solved the issue!!! I have no idea why it was enabled by default”.

As we have noted, [13] both mandates education and does not rely solely on it.

However, as the safety-critical community laments (at least in the U.K. and U.S.A.: cultures do differ here), there is very little training in formal methods for most undergraduates.

## 7 Conclusions

As the media never tire of saying, there are far too many security breaches, and, though they have multiple causes, [12] claims that about 50% of security breaches are caused by coding errors. There appears to be a culture of accepting these, with the U.S. Government investigation [7] into Equifax blaming many factors but not the actual bug, and [13] taking a “necessary but not sufficient” approach to education in secure coding.

**Education** Could certainly do better [57], though there are encouraging signs [58] and useful ideas when it comes to improving informal resources [68]. However, informal resources can be dangerous when it comes to security, and [58] recommends giving *all* students the advice in [69]: “If you pick up a SSL/TLS answer from Stack Overflow, there’s a 70% chance it’s insecure”. More training in formal methods would be welcomed, at least in those cultures where it is lacking.

**Customers/Managers** need to be much more upfront about security requirements [59, 60], and enforce (e.g. by requiring tool support during any CI/CD process, such as [37] describe) at least “middle ground” requirements. In the case of outsourced development, explicit penalty clauses for failing penetration tests should concentrate the developers’ minds.

**C/C++ people** These programmers should be much more aware of techniques for secure coding, such as those described in [46, Appendix F], and the various tools for static analysis.

**Java people** In view of the significance of injection attacks (Number 1 in [10]), programmers should be aware of taint analysis, as in [11].

**JavaScript people** There are some techniques, such as [54], for protecting JavaScript applications, but they are not deployable in the typical JavaScript “dynamic loading web page” environment. Furthermore this environment is basically antithetical to security, as British Airways is learning to the cost of £183M [5].

- 1) Hence the first real challenge of JavaScript lies with the tool makers: there are, as far as the author knows, no JavaScript verifiers in existence, and no page-bundler that checks for version drift, or does incremental verification (which might be comparatively cheap, as in [37]).
- 2) An alternative approach might be to change the JavaScript model. This is advocated in [70], based on their analysis of what third-party scripts do in the wild. This is not a completely radical idea: Google is testing its **TrustedTypes** feature [56], with the motivation “The DOM API is insecure by default and requires special treatment to prevent XSS”.

**Empirical Research** There is not much analysis of the efficacy of various techniques in security programming. [71] compares various techniques, and states the following.

Based on our case study [of two large programs], the most efficient vulnerability discovery technique is automated penetration testing. Static analysis finds more vulnerabilities but the time it takes to classify false positives makes it less efficient than automated testing.

This assumes that “false positives” are acceptable, a debatable point of view. It would be good to have more such research.

**Tool developers** There is a lack of tools (or at least a lack of awareness of tools) that can be neatly integrated into a security programming toolchain the way such tools are integrated in safety-critical toolchains [37].

## Acknowledgements

A predecessor of this paper was given at the 2019 Working Formal Methods Symposium (FROM2019) in Timisoara, Romania. The authors are grateful to the referees and audiences of FROM2019 and FMFun2019 for useful comments. The first author is grateful to the Fulbright Programme for a Cybersecurity Scholarship at New York University in 2017, and to many correspondents and discussions, notably Alastair Irons, Tom Prickett and Tim French.

## References

1. Jacquél, M., Berkani, K., Delahaye, D., Dubois, C.: Verifying B proof rules using deep embedding and automated theorem proving. In: International Conference on Software Engineering and Formal Methods, Springer (2011) 253–268

2. Bloomberg: Equifax Hack Lasted for 76 Days, Compromised 148 Million People, Government Report Says. <http://fortune.com/2018/12/10/equifax-hack-last-ed-for-76-days-compromised-148-million-people-government-report-says/> (2018)
3. Irwin, L.: Marriott downgrades severity of 2018 data breach: 383 million customers affected. <https://www.itgovernance.co.uk/blog/marriott-downgrades-severity-of-2018-data-breach-383-million-customers-affected> (2019)
4. Ford, N.: Medical debt collection agency files for bankruptcy protection after data breach. <https://www.itgovernance.co.uk/blog/medical-debt-collection-agency-files-for-bankruptcy-protection-after-data-breach> (2019)
5. The Guardian: BA faces £183m fine over passenger data breach. <https://www.theguardian.com/business/2019/jul/08/ba-fine-customer-data-breach-british-airways> (2019)
6. Royal Society: Progress and research in cybersecurity: Supporting a resilient and trustworthy system for the UK. <http://royalsociety.org/cybersecurity> (2016)
7. United States Government Accountability Office: Actions Taken by Equifax and Federal Agencies in Response to the 2017 Breach. <https://www.gao.gov/assets/700/694158.pdf> (2018)
8. Osborne, C.: Google patches 'awesome' XSS vulnerability in Gmail dynamic email feature. <https://www.zdnet.com/article/google-patches-awesome-xss-vulnerability-in-gmail/> (2019)
9. Lenart, L.: Security Bulletin S2-045. <https://cwiki.apache.org/confluence/display/WW/S2-045> (2017)
10. Open Web Application Security Project (OWASP): The Ten Most Critical Web Application Security Risks. [https://www.owasp.org/index.php/Category:OWASP\\_Top\\_Ten\\_Project\#tab=Main](https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project\#tab=Main) (2017)
11. Livshits, V., Lam, M.: Finding Security Vulnerabilities in Java Applications with Static Analysis. In: Proceedings USENIX Security Symposium. (2005) 271–286
12. McGraw, G.: Software Security — Building Security In. Addison-Wesley (2006)
13. Payment Card Industry Security Standards Council (PCI SSC): Requirements and Security Assessment Procedures Version 3.2.1. [https://www.pcisecuritystandards.org/documents/PCI\\_DSS\\_v3-2-1.pdf](https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-2-1.pdf) (2018)
14. Ponemon Institute: The State of Web Application Firewalls. Ponemon Institute (2019)
15. Krebs, B.: What We Can Learn from the Capital One Hack. <https://krebsonsecurity.com/tag/capital-one-breach/> (2019)
16. Kolochenko, I.: Web Application Firewall: a must-have security control or an outdated technology? <https://www.csoonline.com/article/3032743/web-application-firewall-a-must-have-security-control-or-an-outdated-technology.html> (2016)
17. Barth, B.: No fly-by-night operation: Researchers suspect Magecart group behind British Airways breach. <https://www.scmagazine.com/home/security-news/no-fly-by-night-operation-researchers-suspect-magecart-group-behind-british-airways-breach/> (2018)
18. Klein, G., Elphinstone, K., Heiser, G., Andronick, J., Cock, D., Derrin, P., Elkaduwe, D., Engelhardt, K., Kolanski, R., Norrish, M., Sewell, T.: seL4: Formal verification of an OS kernel. In Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles (2009) 207–220
19. The Guardian: Hacking risk leads to recall of 500,000 pacemakers due to patient death fears. <https://www.theguardian.com/technology/2017/aug/31/hacking-risk-recall-pacemakers-patient-death-fears-fda-firmware-update> (2017)

20. Newman, L.: Hackers Made an App That Kills to Prove a Point. <https://www.wired.com/story/medtronic-insulin-pump-hack-app> (2019)
21. Evans, M., Loftus, P.: Rattled by Cyberattacks, Hospitals Push Device Makers to Improve Security. <https://www.wsj.com/articles/rattled-by-cyberattacks-hospitals-push-device-makers-to-improve-security-11557662400> (2019)
22. Food and Drug Administration: FDA Informs Patients, Providers and Manufacturers About Potential Cybersecurity Vulnerabilities in Certain Medical Devices with Bluetooth Low Energy. <https://www.fda.gov/news-events/press-announcements/fda-informs-patients-providers-and-manufacturers-about-potential-cybersecurity-vulnerabilities-0> (2020)
23. Heiser, G.: What's new in the world of seL4. [https://archive.fosdem.org/2019/schedule/event/world\\_of\\_sel4/](https://archive.fosdem.org/2019/schedule/event/world_of_sel4/) (2019)
24. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. RFC-8446 (2018)
25. Beck, K. *et al.*: The Agile Manifesto. <http://agilemanifesto.org/> (2001)
26. Blodget, H.: Mark Zuckerberg On Innovation. <https://www.businessinsider.com/mark-zuckerberg-innovation-2009-10> (2009)
27. Lane, A.: Security + Agile = FAIL. [https://securosis.com/assets/library/presentations/Security/AgileFAIL\\_OWASP.ppt\\_.pdf](https://securosis.com/assets/library/presentations/Security/AgileFAIL_OWASP.ppt_.pdf) (2018)
28. Bartsch, S.: Practitioners' Perspectives on Security in Agile Development. In International Conference on Availability Reliability and Security (2011) 479–484
29. van der Heijden, A., Broasca, C., Serebrenik, A.: An empirical perspective on security challenges in large-scale agile software development. In: Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement. ESEM '18, New York, NY, USA, ACM (2018) 45:1–45:4
30. Tahaei, M., Vaniea, K.: A Survey on Developer-Centred Security. [https://groups.inf.ed.ac.uk/tulips/papers/A\\_Survey\\_on\\_Developer\\_Centred\\_Security.pdf](https://groups.inf.ed.ac.uk/tulips/papers/A_Survey_on_Developer_Centred_Security.pdf) (2019)
31. Chapman, R.: Industrial experience with Agile in high-integrity software development. In Parsons, M., Anderson, T., eds.: Developing Safe Systems: Proceedings of the Twenty-fourth Safety-critical Systems Symposium, Safety-Critical Systems Club (2016) 143–154
32. O'Connor, T., Enck, W., Reaves, B.: Blinded and Confused: Uncovering Systemic Flaws in Device Telemetry for Smart-Home Internet of Things. In: Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks. (2019) 140–150
33. Wäyrynen, J., Bodén, M., Boström, G.: Security Engineering and eXtreme Programming: an Impossible Marriage? Extreme programming and agile methods-XP/Agile Universe (2004) 117–128
34. Statt, N.: Zuckerberg: 'Move fast and break things' isn't how Facebook operates anymore. <https://www.cnet.com/news/zuckerberg-move-fast-and-break-things-isnt-how-we-operate-anymore/> (2014)
35. Salz, R.: Software engineering and OpenSSL is not an oxymoron (presentation at Real World Cryptography 2017). <https://rwc.iacr.org/2017/Slides/rich.saltz.pdf> (2017)
36. Seggelmann, R., Tuexen, M., Williams, M.: Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) Heartbeat Extension. <https://tools.ietf.org/html/rfc6520> (2012)
37. Brain, M., Schanda, F.: A lightweight technique for distributed and incremental verification. In Joshi, R., Müller, P., Podelski, A., eds.: Verified Software: The-

- ories, Tools, Experiments. Volume 7152 of LNCS., Berlin–Heidelberg–New York, Springer (January 2012) 114–129
38. Chapman, R., Moy, Y.: AdaCore Technologies for Cyber Security. <https://www.adacore.com/books/adacore-tech-for-cyber-security> (2018)
  39. Distefano, D., Fähndrich, M., Logozzo, F., O’Hearn, P.: Scaling static analyses at Facebook. *Communications of the ACM* **62** (2019) 62–70
  40. Vogels, W.: Proving security at scale with automated reasoning. <https://www.allthingsdistributed.com/2019/05/proving-security-at-scale-with-automated-reasoning.html> (2019)
  41. Sadowski, C., Aftandilian, E., Eagle, A., Miller-Cushion, L., Jaspan, C.: Lessons from building static analysis tools at Google. *Commun. ACM* **61**(4) (2018) 58–66
  42. Open Web Application Security Project (OWASP): DefectDojo: OpenSource Application Security Management. <https://www.defectdojo.org> (2019)
  43. Chapman, R.: Development and Formal Verification of Secure Updates for Embedded Systems (slides from Verification 2018). <http://www.testandverification.com/conferences/verification-futures/vf2018/> (2018)
  44. Google (Project Zero): Oday “In the Wild”. <https://googleprojectzero.blogspot.com/p/0day.html> (2019)
  45. Thomas, G.: A proactive approach to more secure code. <https://msrc-blog.microsoft.com/2019/07/16/a-proactive-approach-to-more-secure-code/> (2019)
  46. Centre for the Protection of National Infrastructure: Rail Code of Practice for Security-Informed Safety. CPNI (2019)
  47. ISO/IEC: TS 17961:2013, Information technology — Programming languages, their environments & system software interfaces — C Secure Coding Rules. <https://www.iso.org/standard/61134.html> (2013)
  48. Chong, N., Cook, B., Kallas, K., Khazem, K., Monteiro, F., Schwartz-Narbonne, D., Tasiran, S., Tautschnig, M., Tuttle, M.: Code-Level Model Checking in the Software Development Workflow. To appear in ICSE-SEIP ’20 (2020)
  49. Cavalcanti, A., Miyazawa, A., Wellings, A., Woodcock, J., Zhao, S.: Java in the Safety-Critical Domain. SETSS 2016: Engineering Trustworthy Software Systems (2017) 110–150
  50. Meng, N., Nagy, S., Yao, D., Zhuang, W., Arango Argoty, G.: Secure coding practices in Java: Challenges and vulnerabilities. In 2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE) (2018) 372–383
  51. Google (Chris Povirk): Denial of Service vulnerability for servers that use Guava and deserialize attacker data. <https://groups.google.com/forum/\#!topic/guava-announce/xqWALw4W1vs/discussion> (2018)
  52. Guarnieri, S., Livshits, B.: GATEKEEPER: Mostly Static Enforcement of Security and Reliability Policies for JavaScript Code. In: USENIX Security Symposium. Volume 10. (2009) 76–85
  53. Meyerovich, L., Livshits, B.: Conscript: Specifying and enforcing fine-grained security policies for javascript in the browser. In: 2010 IEEE Symposium on Security and Privacy, IEEE (2010) 481–496
  54. Maffei, S., Mitchell, J., Taly, A.: Isolating JavaScript with Filters, Rewriting, and Wrappers. In: Proceedings ESORICS 2009. (2009) 505–522
  55. Maffei, S., Taly, A.: Language-based isolation of untrusted Javascript. In: Proceedings 22nd IEEE Computer Security Foundations Symposium. (2009) 77–91
  56. Kotowicz, K.: Trusted Types help prevent Cross-Site Scripting. <https://developers.google.com/web/updates/2019/02/trusted-types> (2019)

57. Cable, J.: Every Computer Science Degree Should Require a Course in Cybersecurity. <https://hbr.org/2019/08/every-computer-science-degree-should-require-a-course-in-cybersecurity> (2019)
58. Crick, T., Davenport, J., Irons, A., Prickett, T.: A UK Case Study on Cybersecurity Education and Accreditation. Proc. FIE 2019 (2019)
59. Naiakshina, A., Danilova, A., Tiefenau, C., Smith, M.: Deception Task Design in Developer Password Studies: Exploring a Student Sample. In Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018). USENIX Association (2018) 297–313
60. Naiakshina, A., Danilova, A., Gerlitz, E., von Zezschwitz, E., Smith, M.: "If you want, I can store the encrypted password": A Password-Storage Field Study with Freelance Developers. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, ACM (2019) 140:1–140:12
61. Cimpanu, C.: 7-Eleven Japanese customers lose \$500,000 due to mobile app flaw. <https://www.zdnet.com/article/7-eleven-japanese-customers-lose-500000-due-to-mobile-app-flaw/> (2019)
62. Biscoe, C.: MyFitnessPal data breach: 150 million app users affected. <https://www.itgovernance.co.uk/blog/myfitnesspal-data-breach-150-million-app-users-affected/> (2018)
63. Blackmon, A.: Macy's hit by data breach. <https://eu.freep.com/story/money/business/2018/07/06/macys-data-breach-online/763074002/> (2018)
64. Inbenta (CEO): Inbenta and the Ticketmaster Data Breach. <http://web.archive.org/web/20181121184620/> (2018)
65. Taylor, C., Sakharkar, S.: ');DROP TABLE textbooks;--: An Argument for SQL Injection Coverage in Database Textbooks. In Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19). ACM (2019) 191–197
66. Fischer, F., Böttinger, K., Xiao, H., Stransky, C., Acar, Y., Backes, M., Fahl, S.: Stack Overflow Considered Harmful? The Impact of Copy&Paste on Android Application Security. 38th IEEE Symposium on Security and Privacy (SP) (2017) 121–136
67. Open Web Application Security Project (OWASP): The Ten Most Critical Web Application Security Risks. [https://www.owasp.org/images/f/f8/OWASP\\_Top\\_10\\_-\\_2013.pdf](https://www.owasp.org/images/f/f8/OWASP_Top_10_-_2013.pdf) (2013)
68. Fischer, F., Xiao, H., Kao, C.Y., Stachelscheid, Y., Johnson, B., Razar, D., Fawkesley, P., Buckley, N., Böttinger, K., Muntean, P., Grossklags, J.: Stack Overflow Considered Helpful! Deep Learning Security Nudges Towards Stronger Cryptography. 28th USENIX Security Symposium (USENIX Security 19) (2019)
69. Chen, M., Fischer, F., Meng, N., Wang, X., Grossklags, J.: How Reliable is the Crowdsourced Knowledge of Security Implementation? <https://arxiv.org/abs/1901.01327> (2019)
70. Zhang, M., Meng, W., Lee, S., Lee, B., Xing, X.: All Your Clicks Belong to Me: Investigating Click Interception on the Web. <https://www.microsoft.com/en-us/research/uploads/prod/2019/03/zhang-observer.pdf> (2019)
71. Austin, A., Williams, L.: One technique is not enough: A comparison of vulnerability discovery techniques. In: Proceedings 2011 International Symposium on Empirical Software Engineering and Measurement. (2011) 97–106