Stanford Blockchain Protocol Analysis and Security Engineering 2017

Introduction & Overview

Byron Gibson

Contents

Introduction & Overview

Thank you all for joining us today. I'll give a brief introduction to the conference, an overview of some of the presentation themes, and conclude with a few thank you's.

Security and information assurance are critical in fiduciary systems. An historic analogy is how the invention of double-entry book-keeping enabled the corporation to exist. Previously such organizations could not scale b/c they could not trust their bookkeeping. But double-entry booking-keeping solved that trust problem and facilitated the creation of economic and social value not previously possible. This is why information security has a direct relationship to fiduciary utility and value. If the security is not robust, there is a cap on the usability and value of the system.

While some aspects of blockchain technology can be contentious, the objective of this conference is to focus on common ground we all care about. Namely, exploring the links between blockchain systems engineering and computer science, and their effects on security and systemic failure risk in production. Many of us were originally motivated by a hope that these decentralized cryptographically-assured fiduciary systems could provide more stable, robust, and reliable infrastructure for the world, and thus facilitate the creation of economic and social value not previously possible.

We hope this conference can make a contribution to that end. Our intent is to solicit potentially impactful, early stage papers, presentations and ideas, which may not yet be ready for formal publication, and which could benefit from constructive feedback from a community of peers. We are also interested in similarly impactful updates to previously published work.

Themes: Assurance and Determinism

Several themes emerged as we assembled the program. One is improving certainty, assurance, and determinism in the behavior of these systems in production.

Turing-Incomplete scripting languages are already deterministic in their behavior, but relatively constrained in their capabilities. A number of research approaches seek to extend their capabilities within the constraint of Turing-Incompleteness.

Turing-Complete scripting languages are already versatile in their capabilities, but more vulnerable to disruption and attacks. A number of research approaches seek to reduce attack surface and improve assurance, while maintaining their versatility.

Among permissioned consensus protocol designs, the intent is to replace well-understood and mature infrastructure with new decentralized successors. The current systems may be purpose-built centralized systems, or decentralized but ad hoc and unwieldy ones. Thus a current research objective is to achieve a more optimal combination of assurance, performance, and cost than either kind of incumbent.

Public permissionless consensus protocols also aim to provide an alternative to existing infrastructure. Bitcoin is currently the standard for security and reliability, but perhaps at a cost to performance, utility and energy consumption. There is a “no-free-lunch” aspect to this optimization problem, and the open question is whether Bitcoin’s set of tradeoffs represent a local or absolute maxima. Thus some research aims to demonstrate the former with a novel, superior architecture, while other research anticipates the latter and attempts performance and/or utility improvements within Bitcoin’s security constraints.

Finally, secure hardware for some cryptographic applications has been in successful deployment for some time, but integrating recent advances and applying all of it to blockchain systems carries new challenges.

Themes: Simulation

Another theme is simulation of complex systems. This has been a well-funded area of computer science research for decades, and is now in the early stages of being applied to decentralized protocol research. While it remains a hard scientific and epistemological problem, it may be fairly tractable when investigating well-specified decentralized protocols, and thus able to provide useful avenues for both research and teaching.

The ability to explore the performance and security assertions of any given protocol in a large-scale but fully controlled environment may fill a useful niche between testing in a small-scale but fully controlled intranet environment, and systems testing in a large-scale but partially-controlled public testnet environment. The ability to control for specific variables at large scale may become a valuable research tool.

Additionally, simulation could provide a useful and fun pedagogical approach for university courses. For example, assigning students the tasks of developing, deploying, running, and even attacking any given blockchain system in a large-scale simulator could provide a holistic hands-on introduction to these systems, as well as help cultivate the much-needed adversarial mindset among future engineers.

Finally, simulators may be useful for comparative studies, penetration testing, and red-teaming, without putting any real value at risk.

These themes are not exhaustive or perfectly representative of all the interesting work related to this conference, but are among the more actively researched at this time.

Conclusion

In closing, I would like to thank the Stanford Cyber Initiative, directed by Allison Berke, and the Stanford Computer Security Lab, co-directed by Dan Boneh, for co-hosting the

conference. I would also like the thank our Program Committee for their early support. Without them it would not have been possible. Special thanks to Elizabeth Stark and Joseph Poon from Lightning Labs, who also helped a great deal.

Thank you also to everyone joining us, especially from out of town and overseas, and to everyone who submitted a presentation. We received more good submissions than we could include in a two day program, are grateful to everyone who submitted, and hope to keep in touch for next time.

Finally, this is a topic many people care deeply about, and there are not many conferences exclusively dedicated to it. Another new one is the IEEE Security and Privacy on the Blockchain Workshop (IEEE S&B), scheduled for April 29th in Paris. We'll send the link to its website in a conference followup email (http://prosecco.gforge.inria.fr/ieee-blockchain2016/).

Thank you, and now on with the program.