Systems, Networking, and Cybersecurity Ph.D. Qualifier Exam
Spring 2019

The following questions relate to the papers published in the reading list on the Spring 2020 qualifier webpage (http://people.cs.vt.edu/~mdhicks2/qual/). For full citations, please see that reading list. Before starting, read and understand the following guide provided by Virginia Tech: Avoiding Plagiarism: A Guide For Graduate Students at Virginia Tech (https://graduateschool.vt.edu/content/dam/graduateschool_vt_edu/graduate-honor-system/avoiding-plagiarism-short-guide.pdf). In your answers, you must avoid unattributed direct quotations and paraphrases and use proper documentation of all sources you use. This requires that you include a bibliography in your response. **Failure to follow these guidelines represents a violation of Virginia Tech's Honor Code and will result in a score of 0.**

Stellar

The following questions relate to the paper on Stellar by Lokhava et al.

a) Provide a brief synopsis of the paper.
b) The paper discusses "issuer-enforced finality" as a key goal. Explain why traditional blockchains exploiting Nakamoto-style consensus fail to provide this goal!
c) What is an example of a "statement" in FBA in the context of Stellar's application to allow currency transfers?
d) What does the author's "Internet Hypothesis" entail?
Discuss (a) the practicality, and (b) threats and risks to the hypothesis.

Algorand

The following questions relate to the paper on Algorand by Gilad et al.

a) Provide a brief synopsis of the paper.
b) What purpose does Algorand's use of "cryptographic sortition" have?
c) Compare Algorand to the original Bitcoin protocol and briefly compare and contrast the two. In your opinion, what is the most important advantage one has over the other? (Consider both directions.)
d) Compare Algorand to Stellar and briefly compare and contrast the two. In your opinion, what is the most important advantage one has over the other? (Consider again both directions.)

Atom

The following questions relate to the paper on Atom by Kwon et al.

a) Provide a brief synopsis of the paper.
b) What are the exact anonymity guarantees provided by Atom?
c) For what purpose does Atom use trap messages?
d) Describe Atom's adversary/threat model and provide your personal assessment of its realism.

Facebook Data Broker

The following questions relate to the paper "Privacy Risks with Facebook's PII-based Targeting: Auditing a Data Broker's Advertising Interface"

a) Briefly describe the privacy attacks presented in this paper. For each attack, identify any aspects of the Facebook advertising interface that enable the attack.
b) For each attack, describe the cost of the attack, and scenarios where it can fail (without considering the specific defenses discussed in the paper). After considering the cost and the failure scenarios, do you think these attacks are still a concern?
c) If Facebook has a perfect Sybil detection system, do you think defenses based on anomaly detection and/or rate-limiting of queries would work? If yes, discuss how you would design an anomaly detection scheme to defend against the proposed attacks? How would the design of your scheme change if Facebook has no existing Sybil detection system, i.e., how would you make your scheme resilient to Sybil attacks?
d) What are the non-privacy threats facing online advertising? From your list, choose any single threat, and provide a detailed explanation of the attack, and associated defenses based on prior work. If there are no existing defenses, suggest your own. For each defense, also identify any limitations.

AdVersarial

The following questions relate to the paper "AdVersarial: Perceptual Ad Blocking meets Adversarial Machine Learning"

a) What is perceptual ad blocking? What are the strengths of the approach when compared to other ad blocking schemes?
b) The paper discusses a perceptual ad blocking scheme called "Percival". Describe the design of Percival. Identify any challenges with training and maintaining a system like Percival. Are false positives (non-ad images considered as ads) a concern? If yes, how would you handle false positives?
c) How would you go about interpreting a given classification decision of Percival, i.e., what aspects of an image led to the classification? Do you think model interpretation schemes can help us better understand false negatives and false positives? Explain.
d) In the paper, adversarial examples are mentioned as a threat against perceptual ad blockers that operate on images. What are adversarial examples (in the context of attacking a visual classifier)? How can an adversary craft adversarial samples to fool a system like Percival? Are such attacks realistic?
e) How would you go about modifying a system like Percival to be effective in the presence adversarial example attacks?

ML Privacy

The following questions related to the paper "Privacy Risks of Securing Machine Learning Models against Adversarial Examples"

a) Briefly describe the key ideas in the paper.
b) Discuss prior work on membership inference attacks against deep neural networks. What aspect of neural networks enable these attacks?
c) Can membership inference attacks be thwarted using regularization schemes? Discuss regularization schemes (not discussed in this paper) that may be effective against the proposed attacks in the paper. Based on your answers, do you think membership inference attacks are a real concern?

DIVA

The following questions relate to "DIVA: a reliable substrate for deep submicron microarchitecture design"

a) Compare and contrast hardware reliability and software reliability.
b) Compare and contrast reliability and security.
c) How does Moore's law impact the challenge of functional verification of hardware and how?
d) Why not just use the DIVA core alone?
e) What are the reliability and performance costs of relying too heavily on the DIVA core (e.g., because my high-performance cores is really buggy) for forward progress?

GLIFT

The following questions relate to "Complete Information Flow Tracking from the Gates Up"

a) How does GLIFT affect implicit information flows?
b) What is the limiting factor in dynamic taint tracking?
c) What is the driving observation of this paper that makes their version of taint tracking tractable?
d) What tradeoff does the GLIFT ISA present compared to traditional ISAs?
e) What is the biggest limitation of GLIFT and how would you address it?

Star-CPU

The following questions relate to "Crafting a Usable Microkernel, Processor, and I/O System with Strict and Provable Information Flow Security"

a) Compare and contrast a microkernel and a monolithic kernel.
b) What fundamental design aspect enables the proposed hybrid (i.e., hardware+software) approach to attain 1/4 the hardware overhead, while being more expressive? How does this design decision effect the practicality of the proposed system?
c) How do the authors support several mutually-distrusting entities running on their system?
d) What purpose does limiting execution time serve in the proposed design?