SARR: A Cybersecurity Metrics and Quantification Framework

Shouhuai Xu, PhD

Gallogly Chair Professor in Cybersecurity

Founding Director, Laboratory for Cybersecurity Dynamics

Department of Computer Science

University of Colorado Colorado Springs

https://xu-lab.org/

10/6/2021 @ CAE Forum

Acknowledgements

This research is not possible without

- ❑ my mentors for moral support and philosophic advices
- ❑ my many collaborators and students for their contributions
- ❑ support from funding agencies

Outline

- The Cybersecurity Metrics and Quantification problem
- The SARR Framework
- Status Quo
- Future Research Directions

A Simple, But Ambitious Question

which I have been thinking for years

- We have many terms/concepts/notions/"buzzwords":
  - Security
  - Dependability
  - Survivability
  - Resilience
  - Agility
  - Trustworthiness
  - Privacy
- Q: What is the "structure/relation" between them that can be leveraged to unify them into a single framework?
  - Easy to understand the question, but hard to answer
- Observation: Cannot tackle it without addressing a fundamental problem, which is …

The Cybersecurity Metrics (and Quantification) Problem

- …perhaps does not need introduction other than mentioning that it has been on multiple Hard Problem Lists
  - [US INFOSEC Research Council 2007]
  - [US NST Council 2011]
  - [SoS Lablets 2015]

Example Illustrating the Difficulty:

How to Quantify Residual Vulnerability?

VulPecker [ACSAC'2016]

VulDeePecker [NDSS'2018]

SySeVR [IEEE TDSC 2021]

μVulDeePecker [IEEE TDSC 2020]

VulDeeLocator [IEEE TDSC 2021]

Robust Detector [under review]

……

Quantifying Software Residual Vulnerability (or Susceptibility)

Why Is Cybersecurity Metris So Hard?

[NSF SaTC 2019 PI Meeting, led by Xu and Trivedi]

1. Systems security is about emergent properties (system vs. components)
2. Hard to precisely define what we want
3. Hard to measure well-defined, useful metrics
4. Hard to parameterize/validate models
5. Walls between sub-disciplines (silos)
6. Technical-organizational misaligned objectives
7. Hard to develop metrics that are reproducible
8. Deal with unknown and future (vulnerabilities, attacks)
9. High dimensionality
10. Context-dependence
11. System complexity
12. Hard to completely specify threat models
13. Hard to relate metrics to threat models
14. Hard to relate vulnerability, exploitability & attack metrics
15. Hard to do experiments at scale
16. Hard to translate intuitive metrics to precise ones
17. Hard to get datasets

This talk presents a systematic approach to overcoming these barriers

Outline

- ❑ The Cybersecurity Metrics and Quantification problem
- ❑ The SARR Framework
  - ❖ Inspired by, and integral to, the Cybersecurity Dynamics approach
- ❑ Status Quo
- ❑ Future Research Directions

The Cybersecurity Dynamics Approach

[Xu2014, Xu2019, Xu2020]

A systematic approach to modeling, **quantifying**, and analyzing cybersecurity from a holistic perspective.

- Using graph structures to describe attack-defense interactions.
- Using parameters to capture attack and defense capabilities, human and software vulnerabilities, etc.
- Using evolution of global cybersecurity state to describe the outcome of attacker-defender-user interactions.

How Is It Different from Others?

- Dynamics-centric
  - Paradigm shift: introducing time into (threat) models
  - Time-independent models → Time-dependent models
- Quantification-driven
  - Quantification isn't an add-on feature but built-in
  - Quantification starts with metrics

Mathematical Abstractions at Nutshell

- Using appropriate mathematical representations
  - Network dynamics G(t)
  - Vulnerability dynamics B(t)
  - Attack dynamics A(t): Dynamic threat models
  - Defense dynamics D(t)
  - Security state metrics M = {m_i} : m_i(t) = $\mathcal{F}_i$ (G(t), B(t), A(t), D(t))
- Example application
  - Compare the effectiveness of architectures and/or mechanisms
  - [illegible]

I will not get into any of these technical details, which are indeed involved/challenging but are not the focus of the present talk

Terminology Used in This Talk

- Levels of abstractions are necessary to cope with cybersecurity
  - Networks: broadly defined to include cyberspace, enterprise networks, infrastructure, cyber-physical-human systems
  - Horizontal view: Network vs. Devices (Computers)
  - Vertical view: Network vs. Components (e.g., hardware, software like OS and IDS, data) vs. Building-Blocks (e.g., TLS)
- Design vs. Operation (a huge gap)
  - Design phase: mostly dealing with building-blocks and components, sometimes with rigorous analysis (e.g., crypto)
  - Operation phase: dealing with networks and devices; rigorous analysis is rare

Terminology (cont.)

- Cybersecurity Properties vs. Security Properties
  - Cybersecurity Properties: broadly defined to include security metrics, agility metrics, resilience metrics, and risk metrics
    - To Do: extension to accommodate dependability, survivability, trustworthiness, privacy
  - Security Properties: narrowly defined to correspond to standard C.I.A., authentication, non-repudiation, etc.
- Metric: A function mapping from a set of objects to a set of value with a certain scale (e.g., {0, 1} or [0, 1]) to reflect cybersecurity properties of the objects
  - Cybersecurity Metrics (broader) vs. Security Metrics (narrower)

SARR Overview

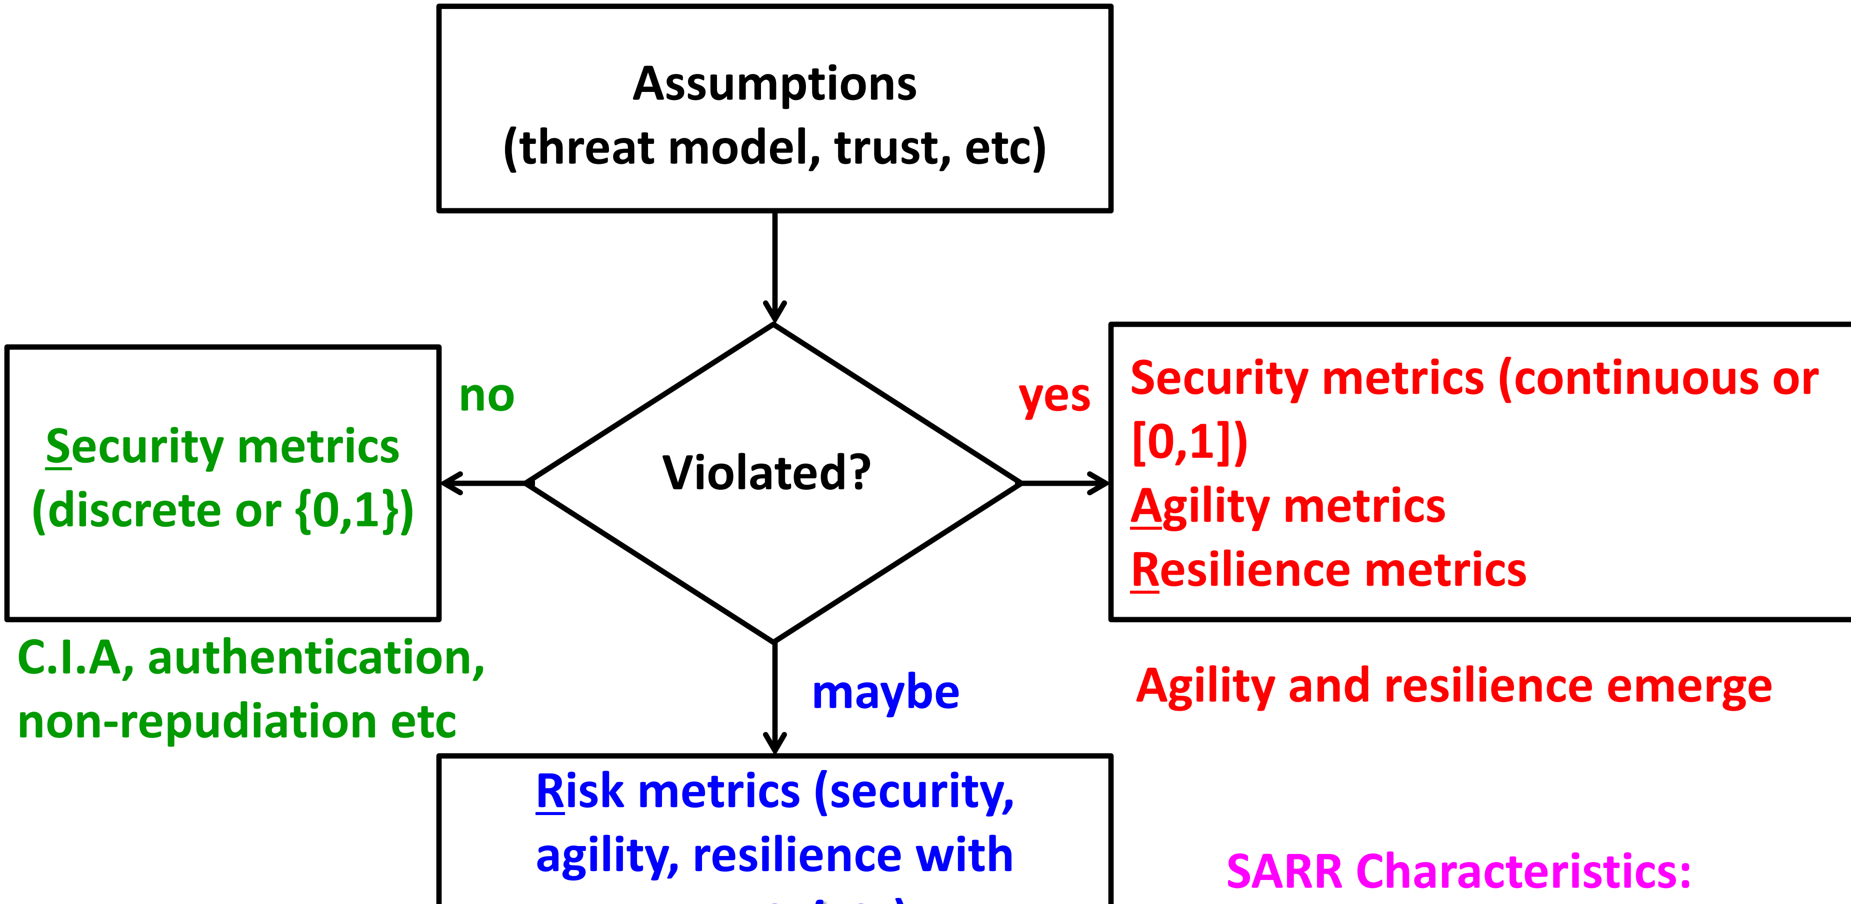

SARR Characteristics:

A next step: Extend it to accommodate dependability (much covered already), survivability (maybe done already), trustworthiness (nothing but conditional probability?), and privacy

Assumptions

- Assumptions associated with the design phase
  - The ones made in the system model, such as: the environment, the communication channel (e.g., private channel vs. authenticated private channel)
  - The ones made in the threat model, such as: chosen-ciphertext attack, adversarial example attack
  - The ones made regarding trust, such as: semi-honest participants
- Assumptions associated with the operation phase
  - The ones “revising or amending” threat model, such as: side-channel capable or not, bounded compromises (1/3 in BFT)

Metrics When Assumptions Not Violated

- Security properties are often discrete or binary, namely {0,1}
  - Often (rigorously) analyzed by designers
  - Often dealing with building-blocks and sometime components, rarely dealing with networks and devices; the latter is often left as "practitioner's problem"
- Metrics associated with the design phase
  - Properties: C.I.A., authentication, non-repudiation, etc.
  - Need precise statement: "property of p holds in what system model against what attacks"
- Metrics associated with the operation phase
  - Service response time and throughput, etc

Metrics When Assumptions Violated

- To what degrees assumptions are violated (with certainty)?
- To what degrees security properties are compromised?
- Agility and resilience come to play
  - Agility: how fast defender reacts to changes (e.g., detecting attacks, responding to attacks)
  - Resilience: degrees of networks/devices/components/building-blocks bouncing back from compromised security properties and violated assumptions; bounceability threshold
- Primarily applicable to the operation phase but having not been systematically investigated: security-by-design (investigated more) vs. agility-by-design vs. resilience-by-design (little understood)

Metrics When Assumptions May Be Violated

- Somewhere in between the two ends of the two spectrum mentioned above: assumptions certainly not violated vs. violated
- Uncertainty comes to play
- What is degree of certainty assumptions are violated?
- What is degree of certainty security properties are compromised?
- What is degree of certainty an alert/anomaly is an attack?
- What is degree of certainty software contains 0-day vulnerability?

Observation 1: Uncertainty is inherent to cybersecurity, so is risk.

Outline

- The Cybersecurity Metrics and Quantification problem
- The SARR Framework
  - Inspired by, and integral to, the Cybersecurity Dynamics approach
- Status Quo
- Future Research Directions

Assumptions

- Often made informally (exception: crypto)
- Often made implicitly
  - E.g., secrecy of cryptographic key → "cryptographic security property ≠ cybersecurity property" → putting trustworthiness of digital signatures or non-repudiation in question
- May be inadequate / incomplete
  - E.g., chosen-plaintext attack → chosen-ciphertext attack
  - E.g., assuming away side-channel attacks → considering them

Observation 2: We must explicitly and precisely articulate assumptions

Security Metrics

via the Cybersecurity Dynamics approach [Pendleton2016]

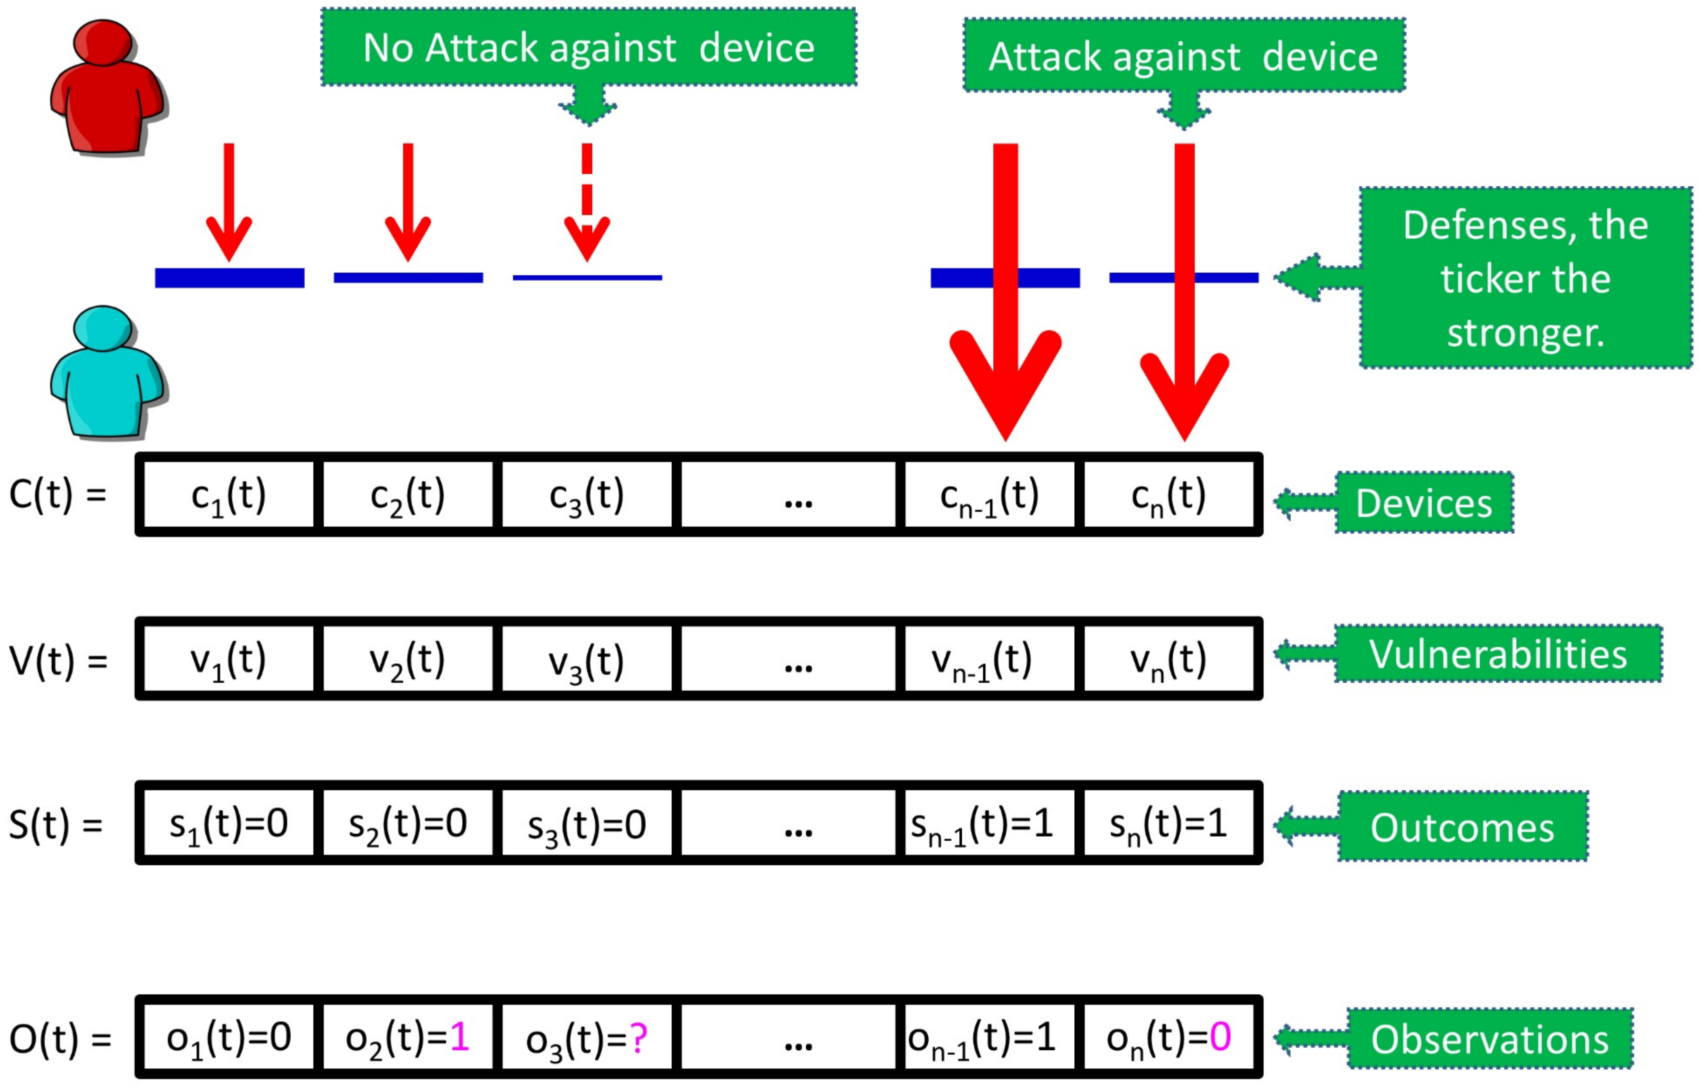

Security Metrics

via the Cybersecurity Dynamics approach [Pendleton2016]

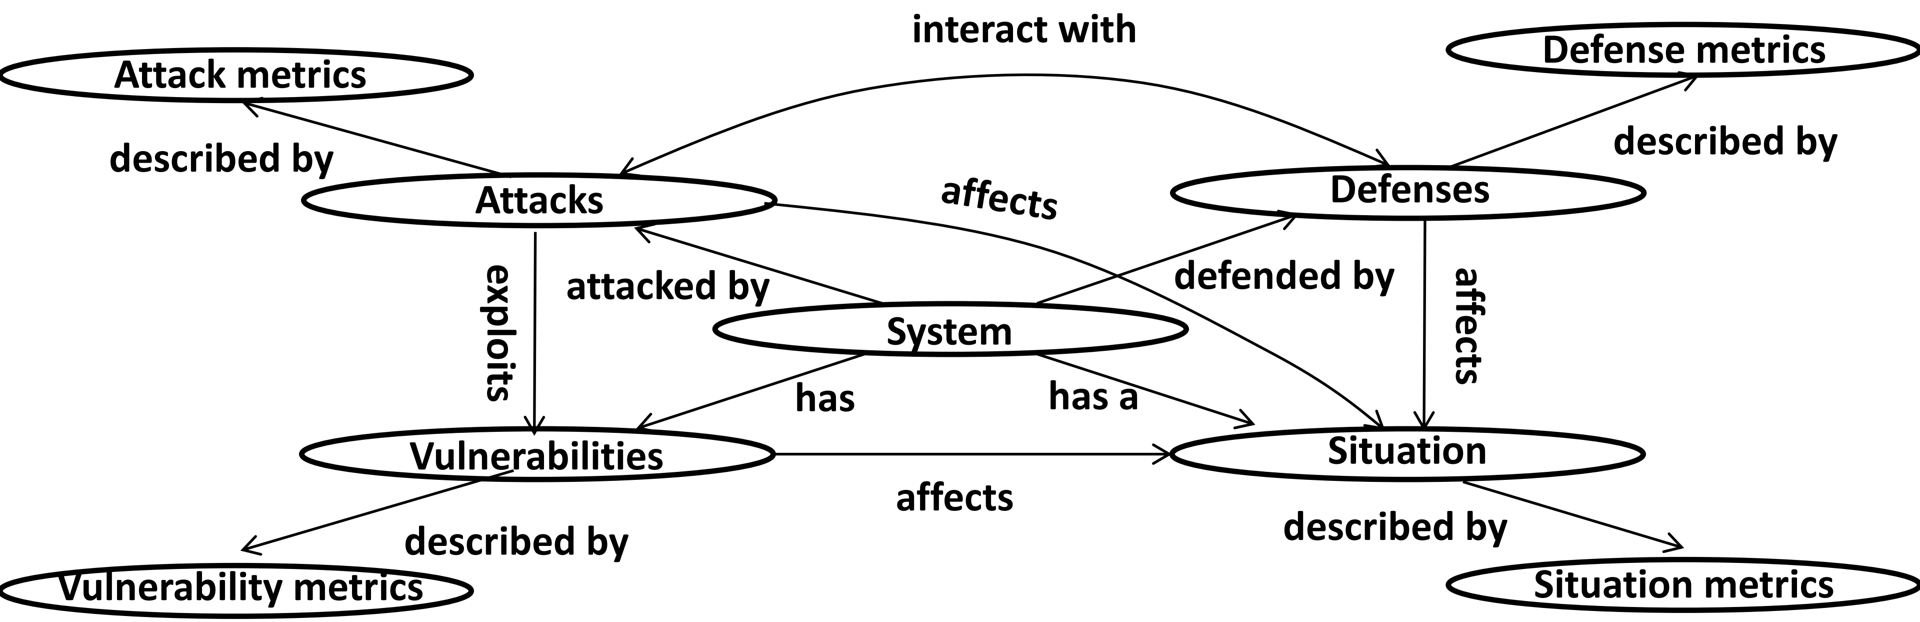

Security metrics = vulnerability metrics ∪ defense metrics ∪ attack metrics ∪ situation metrics

Security Metrics

via the Cybersecurity Dynamics approach [Pendleton2016]

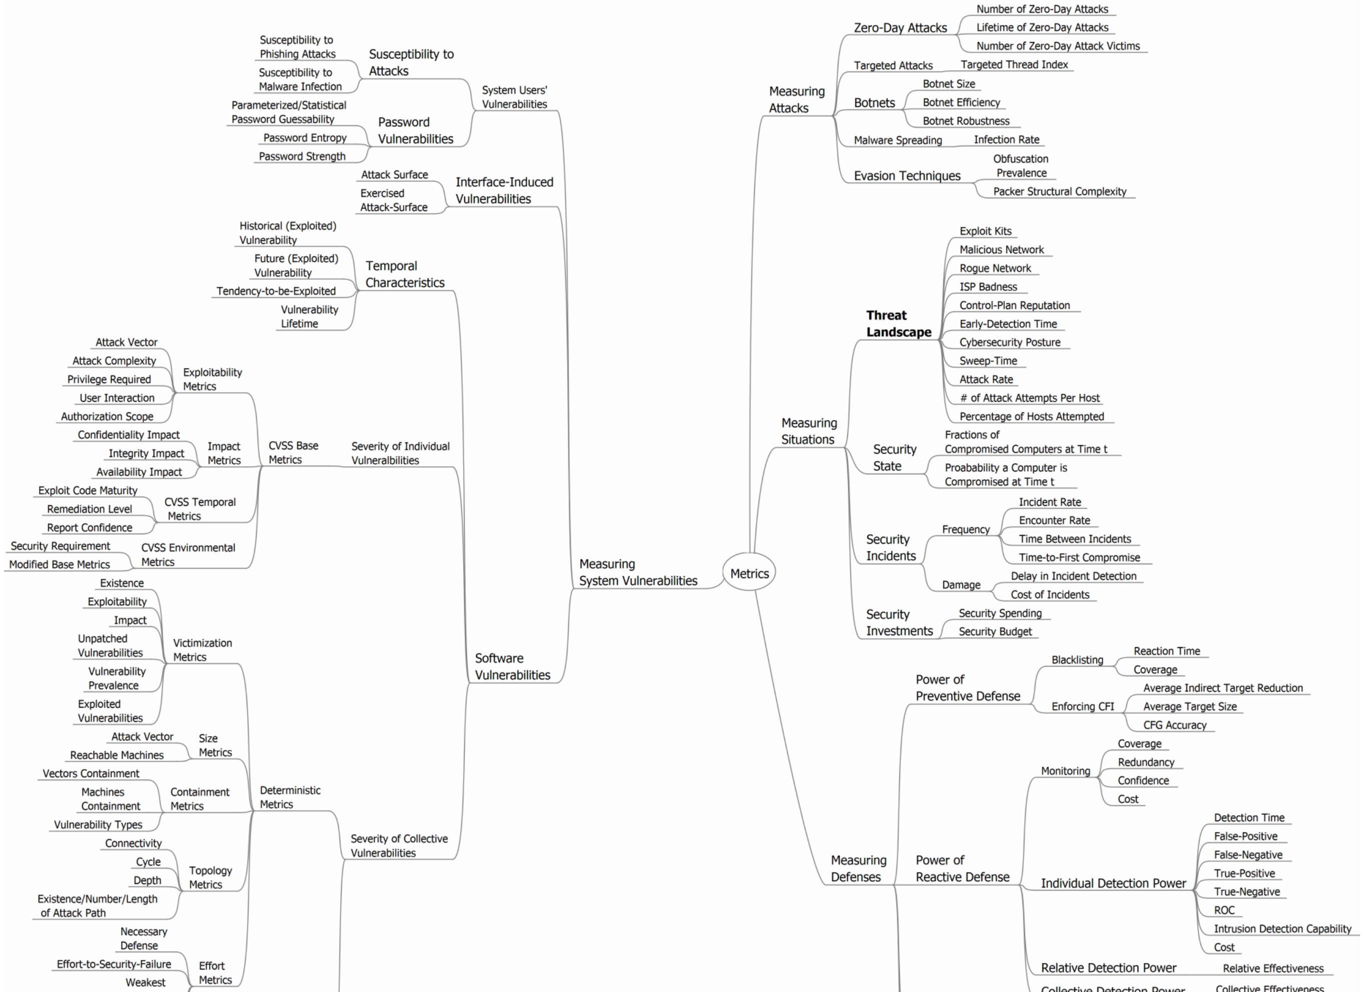

Observation 3: Our understanding of what should be measured is superficial (despite the many metrics)

Gaps in Cybersecurity Metrics

via the Cybersecurity Dynamics approach [Pendleton2016]

What we can do now

- ❑Quantify building-block properties
- ❑What can be measured
- ❑No metrics curriculum
- ❑"1 + 1 + 1 = ?" in the current partnership?
- ❑Most security papers offer no metrics
- ❑Ad hoc definitions of metrics
- ❑Uncertainty largely ignored
- ❑No research community

What need to be done

- ❑ Quantify holistic system properties
- ❑ What must be measured
- ❑ Metrics curriculum
- ❑ Government & industry & academia: 1+1+1>3
- ❑ Each security paper has clearly defined metrics
- ❑ Clear understanding of metrics (e.g., additivity?)
- ❑ Theory of uncertainty quantification
- ❑ A research community

Agility Metrics
via the Cybersecurity Dynamics approach [Mireles2019]

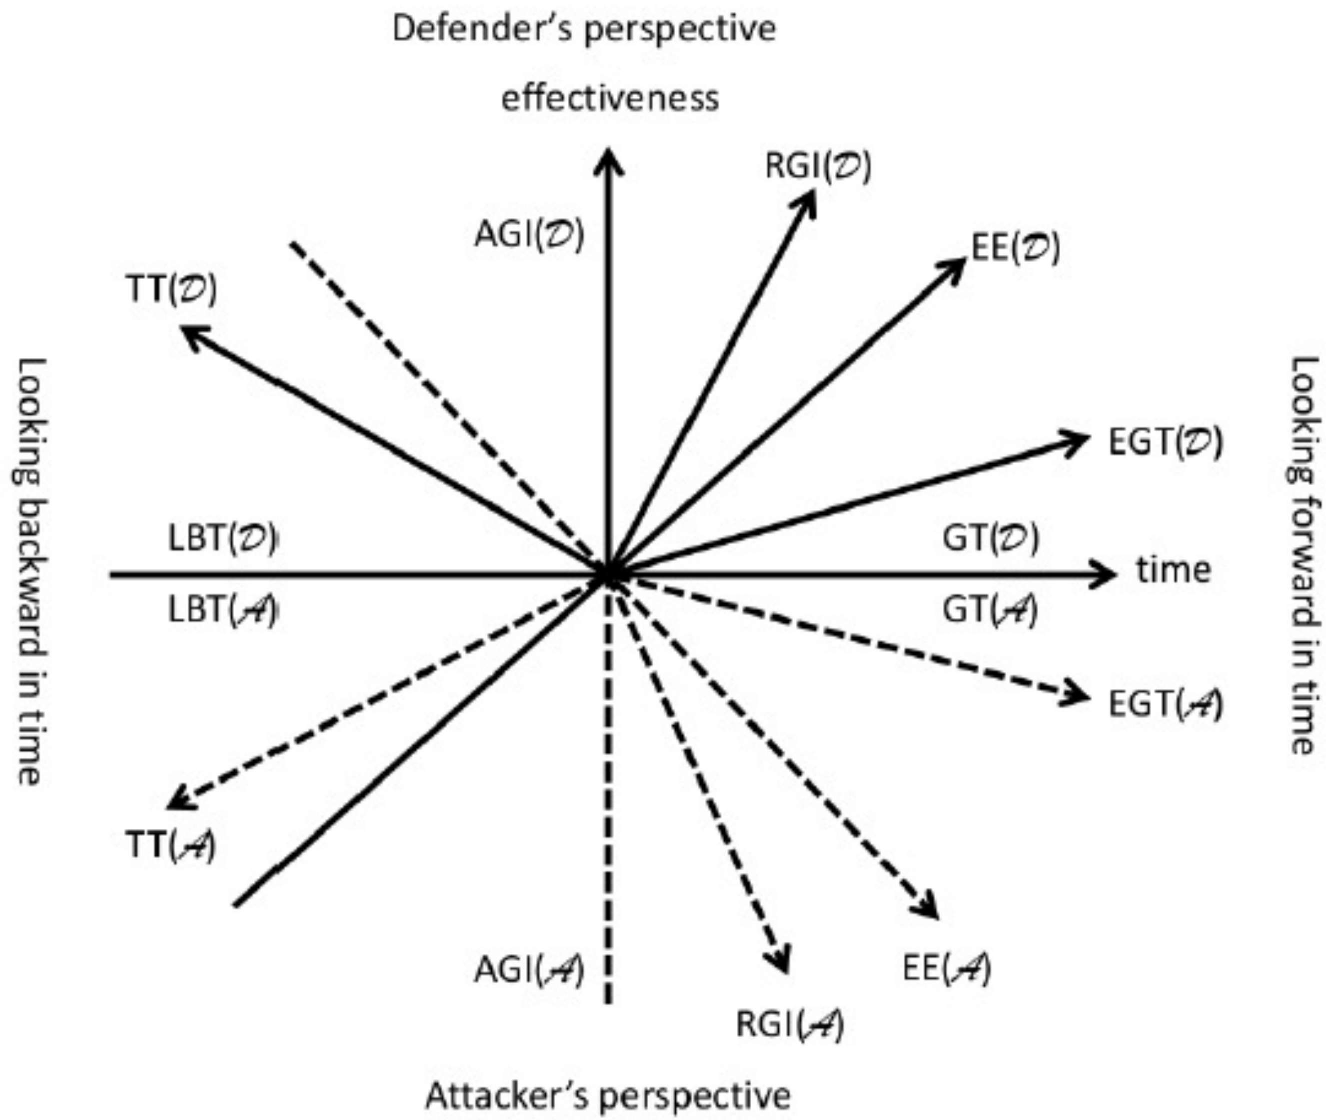

Agility Metrics

via the Cybersecurity Dynamics approach [Mireles2019]

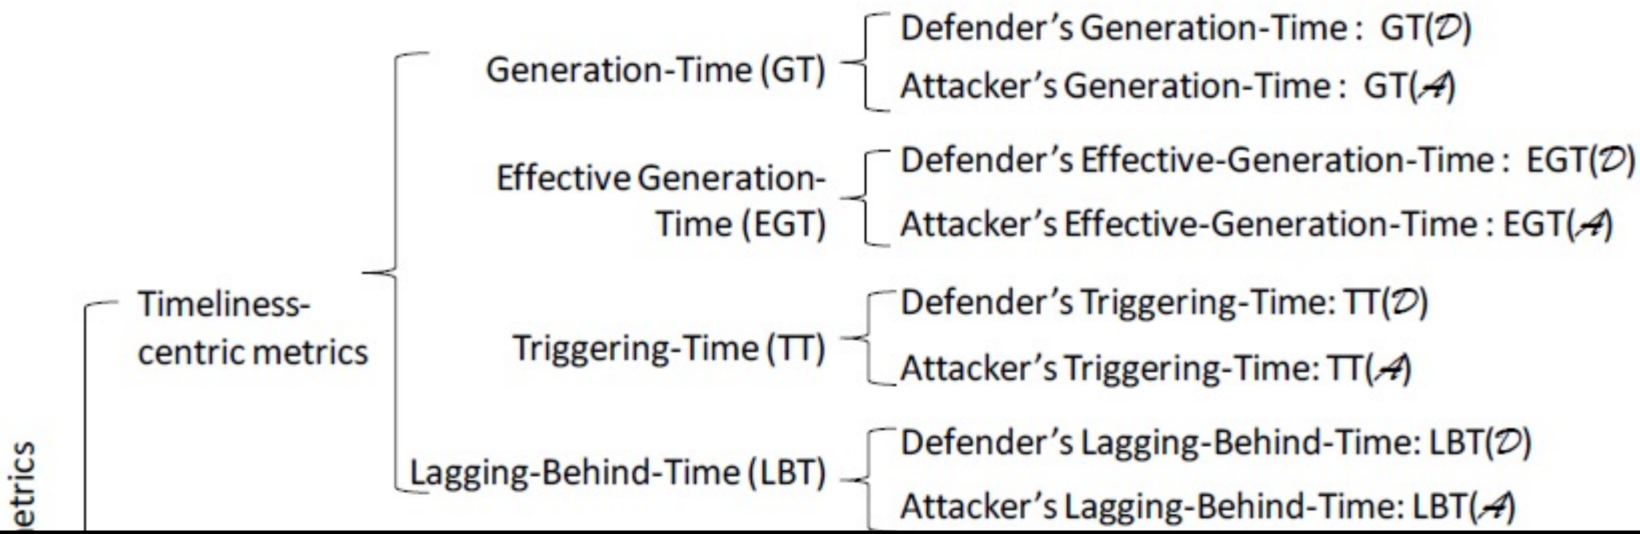

Insights drawn from case study (by applying agility metrics):

- ❑ **Snort is responsive to attacks by timely evolving its defense, but attacks also evolve (i.e., arm race in attack-defense interactions)**
- ❑ **Snort has a lower agility in response to manual attacks than automatic attacks**

Observation 4: Our understanding of agility metrics is superficial

Resilience Metrics [Cho2019]

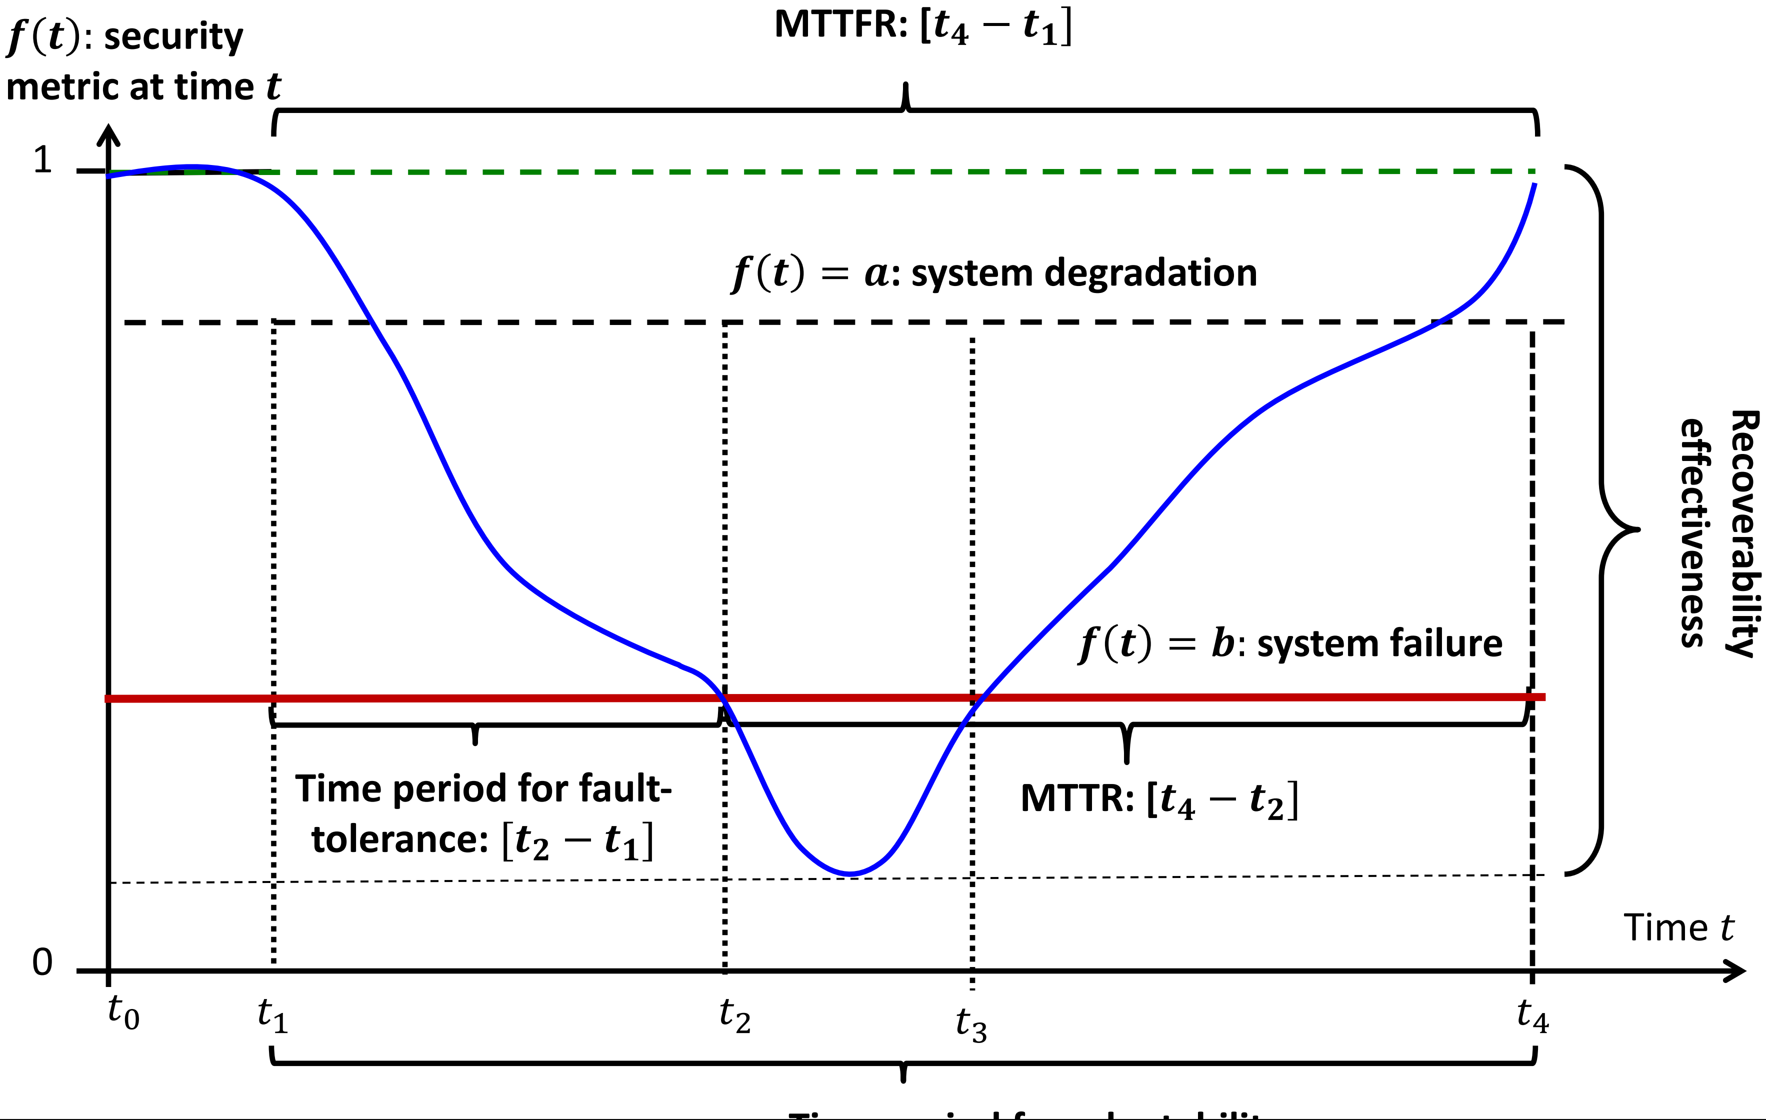

Observation 4: Our understanding of resilience metrics is superficial

Risk Metrics

- Widely used formula (originally proposed to deal with hazards)

  risk = threat × vulnerability × consequence

- Having been "borrowed" to deal with cybersecurity risks, without challenging its applicability
- Not applicable to cybersecurity in general (see references in paper)
  - Do not consider dependence, interdependence, cascading failures, or emergent properties
  - Do not consider the time dimension (or dynamics), by oversimplifying the problem
- The Cybersecurity Dynamics approach aims to overcome them

Observation 5: Our understanding of risk metrics is superficial

Outline

- The Cybersecurity Metrics and Quantification problem
- The SARR Framework
  - Inspired by, and integral to, the Cybersecurity Dynamics approach
- Status Quo
- Future Research Directions

(1) Taming Cybersecurity Assumptions

- The ideal case
  - Assumptions are stated explicitly and precisely
  - Assumptions are independent of each other
  - Assumptions made at design phase are satisfied at operation
- Hard to achieve, but have to do it!
- Alternatives:
  - Characterizing (inter)dependence between assumptions
  - Example: the authenticated private channel assumption depends on the assumption that communication end parties are not compromised, which may further depend on other assumptions (and may even lead to circular assumptions)

(2) Bridging Design vs. Operation Gaps

- The gaps are incurred by
  - Multiple levels of abstractions: design often deals with building-blocks and components (low levels of abstractions) vs. operation often deals with networks and devices (high levels)
    - Speak different languages: “English vs. French” problem
  - Designers assume assumptions will not be violated, but defenders deal with the situations where they are violated
  - Designers may not tell (or care) the operation-phase implications of assumptions made at the design phase

(3) Identifying Metrics That Must Be Measured

- We don't know what metrics we must measure (despite efforts)
- Maybe a useful approach, using medical science as analogy
  - Metrics for building-block or "cell" level cybersecurity properties → "tissue" level cybersecurity properties → "organ" level cybersecurity properties → "human body" level cybersecurity properties
- Emergent property would be reflected by metrics

(4) How Can We Tell Good vs. Poor Metrics?

- Defining metrics are not hard; defining “good” metrics are
  - Analogy: good security definition vs. poor security definition in cryptography
- But what are “good” metrics? According to what criteria?
- How to approach the problem?
- Conduct case studies for some killer applications (e.g., cyber defense command-and-control, quantitative cybersecurity management); need quality data for case studies

(5) Fostering a Research Community

- SciSec and HotSoS are perfect homes for this community
- "Grass roots" approach: Each paper with explicitly and precisely defined assumptions, metrics, and quantitative statements on the progress made by the paper (e.g., security improvement)
  - Rather than: a new attack defeats a defense, or a new defense defeats an attack, without quantitative statements

(6) Developing a Science of Measurement

- Given well-defined cybersecurity metrics, one would think their measurement would be trivial
- May be true sometimes
- But can be extremely challenging → need principled solutions
  - E.g., inferring cybersecurity metrics in the absence of ground-truth
  - Analogy: how is light speed or gravity or time precisely measured in Physics?

Takeaway

Cybersecurity Metrics and Quantification is one of the most fundamental problems to work on (in any context)!

- ❖ **Substantial progresses can be made**
- ❖ **Cybersecurity Dynamics is promising approach**
- ❖ **What are the other approaches?**

❑ **I plan to create materials for "Cybersecurity Metrics" course**

❑ **Yes, we know how hard the problem is, but**

❑ **"Wir müssen wissen, wir werden wissen." ("We must know. We will know.")**

— David Hilbert