Re-architecting tomorrow's internet for "survivability" (a resilience engineering perspective)

David Alderson[1,3] John Allspaw[3] David Woods[2,3]

[1] Naval Postgraduate School, Monterey, CA

[2] The Ohio State University, Columbus, OH

[3] Adaptive Capacity Labs, Brooklyn, NY

NSF Workshop: Towards Re-architecting Today's Internet for Survivability
Evanston, IL | November 28-29, 2023

Part 1:
Survivability in Complex Systems

Survivability & Complex Systems: How Complex Systems Fail

1990 J. Reason & others

Multiple Contributors each necessary only jointly sufficient;
contributors come from multiple system layers
present well before the triggering event

1994 Woods et al. *Behind Human Error*
1997 Reason *Organizational Accidents*
1998 Cook *How complex systems fail*

- How complex systems fail
- How learning breakdowns after failures

eg Novelty Inequality;
Component Substitution Fallacy

2000 Doyle
2003/05/06 Woods

Failure is due to brittle systems, not limited components, subsystems, or human beings

2011 Woods

How Adaptive Systems Breakdown (3 basic patterns)

RyF → Signature of Complex System :

Surprising sudden collapse
against backdrop of continuous improvement / new capabilities

systems "which are robust to perturbations they were *designed to handle*, yet fragile to *unexpected perturbations* and *design flaws*"* (Carlson and Doyle 2000, p. 2529)

- highly competent when events fall within the envelope of designed-for-uncertainties

- sudden, large failures occur in the face of events that challenge or go beyond the envelope

* Limits given T/O choices are "flaws" only in hindsight

Patterns of Adaptive Breakdown

Getting stuck in outdated models.

the world changes but the system remains stuck in what were previously adaptive strategies.

Working at cross-purposes: behavior that is locally adaptive, but globally maladaptive.

inability to coordinate across roles, units, & echelons as goals conflict.

Decompensation: exhausting capacity to adapt as disturbances/challenges cascade.

breakdown occurs when challenges grow and cascade faster than responses can be decided on and deployed to effect.

Survivability & Complex Systems: How Complex Systems Fail

2000-06 Failure is due to brittle systems, not limited components, subsystems, or human beings

2018 / 2019 Doyle DeSS Architectural principles to overcome risks from
Woods TGE brittleness (N dim T/O space)

Pragmatic steps for organizations now?

Continuous learning from stream of incidents
Patterns in how saturation develops, migrates, spreads
Reciprocity, …

How to use new architectural ideas in working systems?

Polycentric, Perspective shifts, …

How to measure future adaptive capacity?

stories of technology change describe or envision
the new forms of *congestion, cascade & conflict* that arise
when apparent benefits get hijacked

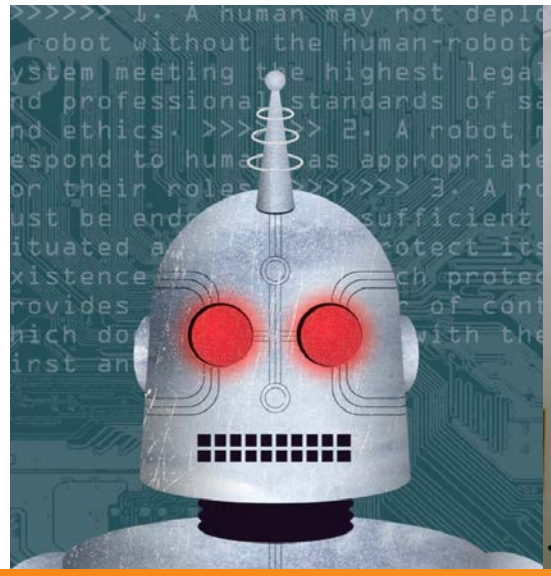

SNAFU CATCHERS

All stories of technology are human stories …

Part 2:
The Internet as Operated

The Internet as a Critical System

It is not sufficient for it to be designed and deployed.

It needs to be ***operated***.

The Internet as a Critical System

You will never have complete knowledge of the system
(components, software, users)

There needs to be as much attention on how its ***operated***
as how it is designed.

The Internet as a Critical System

You will never have complete knowledge of the system
(components, software, users)

You can only ***learn by operating*** it.

The system is always adapting. Can we learn fast enough?

How to guarantee non-survival

Ask (variations of):

- “Why did it break?”
- “Who broke it?”
- “What violations of rules/procedures were made?”

Assume answers to the above will
“prevent it from ever happening again.”


Enjoy
Hindsight Bias

For the people actually responsible for operating the stuff…

- How do they it expect it to work?
- How did it actually work that day?
- Was this a difficult case handled well, or a straightforward case handled poorly?
- What made it so?
- What remains unclear about the event?
- How do others (non-responders, not responsible) understand the event?
- How does their ^^ understanding differ from how responders understand it?

How we imagine incidents happen

- Problems of compliance
- Need to find the root cause
- Can be categorized in a taxonomy, measured, and usefully described with statistics
- Humans are often the problem

How incidents actually happen

- Things are always messy
- Root cause analysis is a fallacy that hides the real problems lurking in system complexity
- Taxonomies often hide rather than reveal; statistics like availability and MTTF are not useful
- Human error as a red herring – (some) people in some roles fill the gaps so that systems work

Survivability as imagined

Survivability as actually happens

Reference:
Hollnagel, Woods & Leveson (2006). Resilience Engineering. Woods et al., Behind Human Error (1994/ 2nd edition 2010). Cook et al. (2009). Minding the Gaps. Quotes from Woods et al. (2021). Patterns in How People Think and Work: Importance of Patterns Discovery for Understanding Complex Adaptive Systems..

Part 3:
Looking Forward

Simulation?

Current strategy:
- Mod & Sim at scale
- Find the gaps
- Fill the gaps

There is no staging environment that is representative of real production systems

Digital twins will not close the gap!

What can we learn from M&S?

Journal of Critical Infrastructure Policy • Volume 2, Number 2 • Fall / Winter 2021

Strategic Perspectives

Progress toward Resilient Infrastructures: Are we falling behind the pace of events and changing threats?

David D. Woods[1] and David L. Alderson[2]

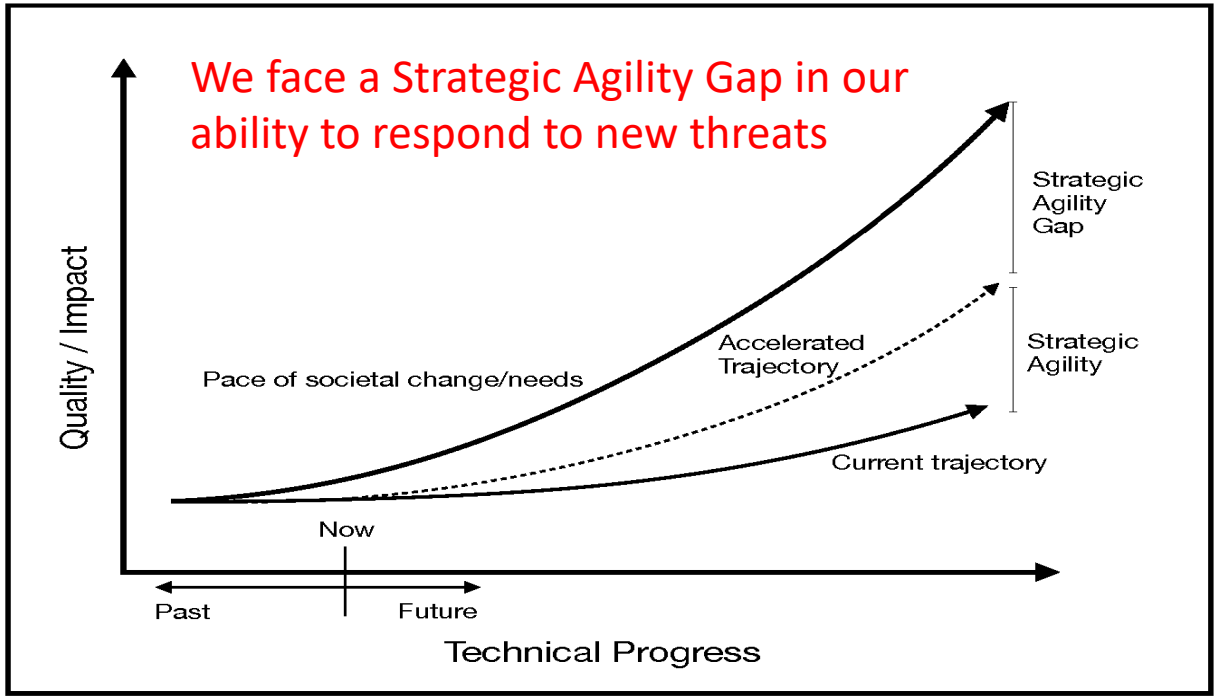

Stress testing?

From Bustamante & Willinger (2023):

> the US Federal Reserve conducts stress tests annually "to assess whether banks are sufficiently capitalized to absorb losses during stressful conditions while meeting obligations to creditors and counterparts and continuing to be able to lend to households and businesses."

https://www.federalreserve.gov/supervisionreg/stress-tests-capital-planning.htm

What can you learn from a stress test?

The Internet as a Critical System

What is ***"critical"*** is going to be dynamic…

Issue #1:
Technology transitions without the requisite attention to safety critical design

Issue #2:
Different stakeholders can treat the same event at different points along this progression!

Issue #3:
The De-Skilling Paradox
(degradation in skills that are no longer practiced regularly)

The Slide to Criticality

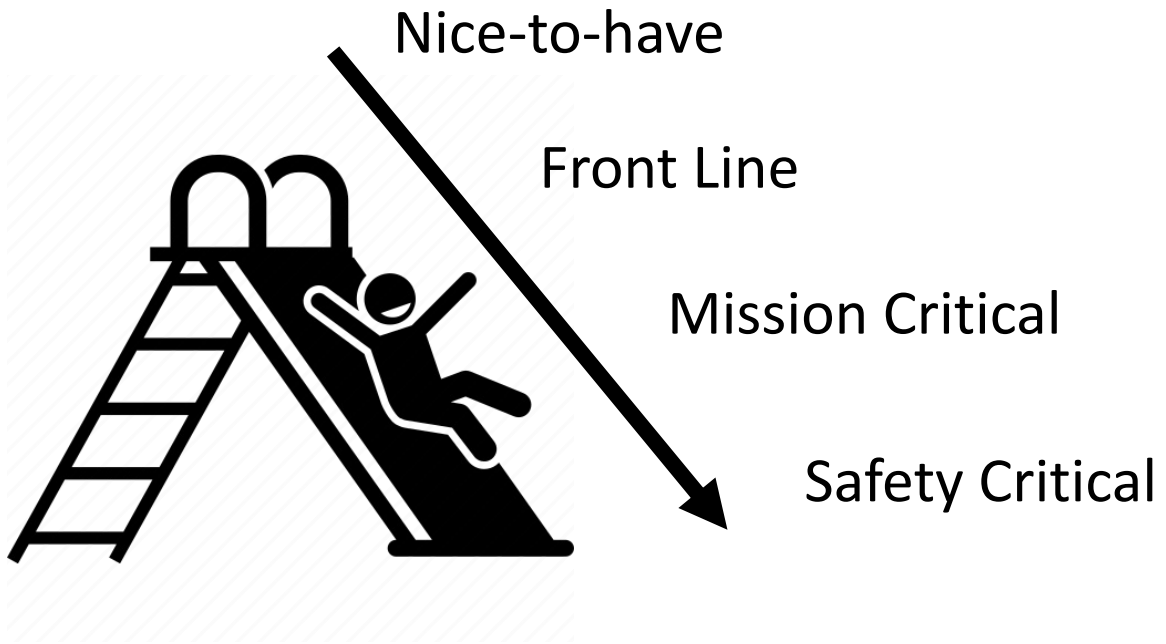

New lines of inquiry

<u>Adaptive Capacity</u>

A system's capacity to adapt to challenges ahead,
when the exact challenge to be handled
***cannot* be specified completely in advance.**

How do we measure adaptive capacity?

New lines of inquiry

Any single perspective simultaneously reveals and obscures.

The way around this is to be able to ***shift perspectives***

How can internet architecture provide
the means to shift perspective?

New lines of inquiry

Governance

A significant but partial rearchitecting of the Internet would require a rearchitecting of governance