

# Networking Panel Summary

- **Panelists**
  - **Jeannie Albrecht**
    - Williams College, PlanetLab testbed
  - **Ivan Seskar**
    - Rutgers University, ORBIT testbed
  - **Steven Schwab**
    - Cobham Analytic Solutions, DETER testbed
  - **Eric Eide**
    - University of Utah, Emulab testbed

# PlanetLab

- PlanetLab: 1000+ *shared* Linux machines with “simplistic” mechanisms for isolation and allocation
- Archival and repeatability requires ability to capture and configure properties of testbed
- Experiment archival is possible on PlanetLab
  - Monitoring and measurement services exist
  - Make it possible to capture live conditions
- Experiment repeatability is not possible
- Possible solution: Use measurements from “live deployment” on PlanetLab to configure Emulab virtual network

# ORBIT

- ORBIT: 400 wireless nodes with support for topology configuration
- Monitoring and measurement tools for archival are built into testbed infrastructure
- Full repeatability/replayability is complicated by characteristics of hardware (very hard to virtualize hardware!)
- Short term repeatability is possible, but long term is not due to hardware degradation
- Difficult to fully capture “state of devices” in testbed, which is necessary for full repeatability

# DETER

- DETER: ~500 nodes running Emulab software with support for various OSes; designed for cybersecurity experiments
- Hardware matters! (Different network cards use different chipsets that have variable performance)
- Low level hardware devices get replaced in testbed and we can't control or prevent that
- Full reproducibility/repeatability is very hard
- Key insight: don't need full repeatability; just keep track of invariants (i.e., the properties/phenomena that matter) rather than contingencies (e.g., hardware, firmware, etc.)

# Emulab

- Emulab: ~500 nodes capable of performing full network emulation (configurable topology)
- Goal should be replayability, not just repeatability
- May want to replay a slightly modified version of experiment; this is even harder to achieve!
- Long term access to physical resources is not possible; hardware degrades and gets replaced
- “Perfect” replayability is not a reasonable goal
- We should strive for replayability with respect to certain invariants

# Discussion

- Methodology that we use in our experiments is roughly the same as other disciplines
  - Metrics are likely very different
  - Testbed monitoring and anomaly detection/reporting in testbeds is essential
  - We keep traces, but people generally don't care
- Hard to classify invariants and contingencies – depends on experiment and testbed

# Discussion

- Need both “wild” and uncontrollable experimentation combined with controlled experimentation in configurable environments
  - Difference between engineering and science
- Qualitative results matter more over time than quantitative results

# Conclusions

- Experiment archival in networking testbeds is possible (although potentially expensive)
- Experimenters need tools for easily obtaining measurement/monitoring data for archival
- “Perfect” repeatability is not really possible due to hardware degradation, calibration, etc
- The goal should be to provide repeatability with respect to certain metrics (or invariants)
- Should provide repeatable qualitative results