

Notes from Networking Panel

- I. Panelist presentation summaries
 - a. Jeannie
 - i. PlanetLab is a live deployment testbed that allows researchers to run experiments “in the wild”
 - ii. PlanetLab is unpredictable and volatile; many tools exist to cope with unpredictability: CoMon, S³, Gush, SWORD are still online
 - iii. Two requirements for experiment archival and repeatability: live capture of current conditions and ability to configure conditions in testbed
 - iv. PlanetLab has (imperfect) support for live capture, but not for configuration!
 - v. Therefore archival is possible, but repeatability is not
 - vi. Questions: Is this bad? Is repeatability required in ALL testbeds? What is repeatability anyway?
 - vii. Possible “solution”: Use conditions captured from PlanetLab to configure more controllable testbeds (like Emulab)
 - viii. Question from audience: If an experiment is not repeatable, is it even valid?
 - ix. Steve’s response to audience: Important to have both controllable and “uncontrollable” experiments
 - b. Ivan:
 - i. Hard to virtualize wireless characteristics
 - ii. Full replayability/repeatability is not really possible
 - iii. Monitoring/measurement services should be built into testbed infrastructure
 - iv. ORBIT is more configurable than PlanetLab because it is not a shared testbed in the sense that only one experiment runs at any point in time
 - v. Very hard to “repeat” hardware properties, and hardware really matters!
 - vi. Temporal repeatability: Experiments can be repeated in short term, but not in the long term
 - vii. Spatial repeatability is also very hard
 - viii. Problem: Cant’ capture “state of devices” which makes repeatability challenging
 - c. Steve:
 - i. Cybersecurity experiments are different from other networking experiments
 - ii. Low level devices fail and get replaced in testbeds; we can’t control that!
 - iii. For example, chipset used on Ethernet controllers really impacts results

- iv. Therefore full reproducibility is very, very hard
- v. Steve supports notion of repeatable experiments in theory, but in practice it's just not full possible in networks – too many variables!
- vi. You don't need all details – only need invariants
- vii. Publications don't capture full details in software complexity
- viii. Archival is expensive! How much are we willing to pay?
- ix. Need to be able to reuse *parts* of an experiment
- x. Comment from audience: In biology, we use different populations with same characteristics to verify repeatability. How relevant is this to networks?

d. Eric

- i. Goal should be replayability, not just repeatability
- ii. Emulab is unique (compared to other testbeds) in that it allows network topology configuration
- iii. Hard to package/capture semantics and meta data
- iv. OS functions are somewhat non-deterministic, capturing can help, but it is somewhat expensive
- v. May not want to repeat exact same experiment, but rather repeat a slightly modified experiment – this is even harder!
- vi. Long term access to physical resources is not possible. Virtualization helps a little, but not everything can be virtualized
- vii. Perfect replayability is not a reasonable goal; we should strive for replayability with respect to certain invariants (which requires being able to identify what matters in an experiment)

II. Discussion

- a. How are we different from supercomputers with batch schedulers? Methodology and environment?
 - i. Methodology is mostly the same
 - ii. Metrics that we care about are arguably different in most cases
 - iii. We monitor testbed usage and status regularly
 - iv. Report changes or over-subscriptions
 - v. Anomaly detection and reporting is important
- b. How do we identify invariants vs. contingencies?
 - i. Good question! Depends on experiment and testbed
- c. Can you record statistical measurements about testbed?
 - i. We follow a different model from chemistry
 - ii. We do keep traces, but no one cares about them
- d. We need both “wild” deployment and controlled experimentation
 - i. Difference between engineering and science
- e. How relevant is repeatability to “real life”?
 - i. Data doesn't matter as much as high level results
 - ii. Care about qualitative results rather than quantitative results