RPKI Trust Anchor Usage and Cache Consistency

John Kristoff <jtk@depaul.edu>

DePaul University (ops role)
University of Illinois at Chicago (research role)

In collaboration with:

Randy Bush <randy@psg.com>
George Michaelson <ggm@apnic.net>
Thomas C. Schmidt <t.schmidt@haw-hamburg.de>
Matthias Wählisch <m.waelisch@fu-berlin.de>

Overview of RPKI validation

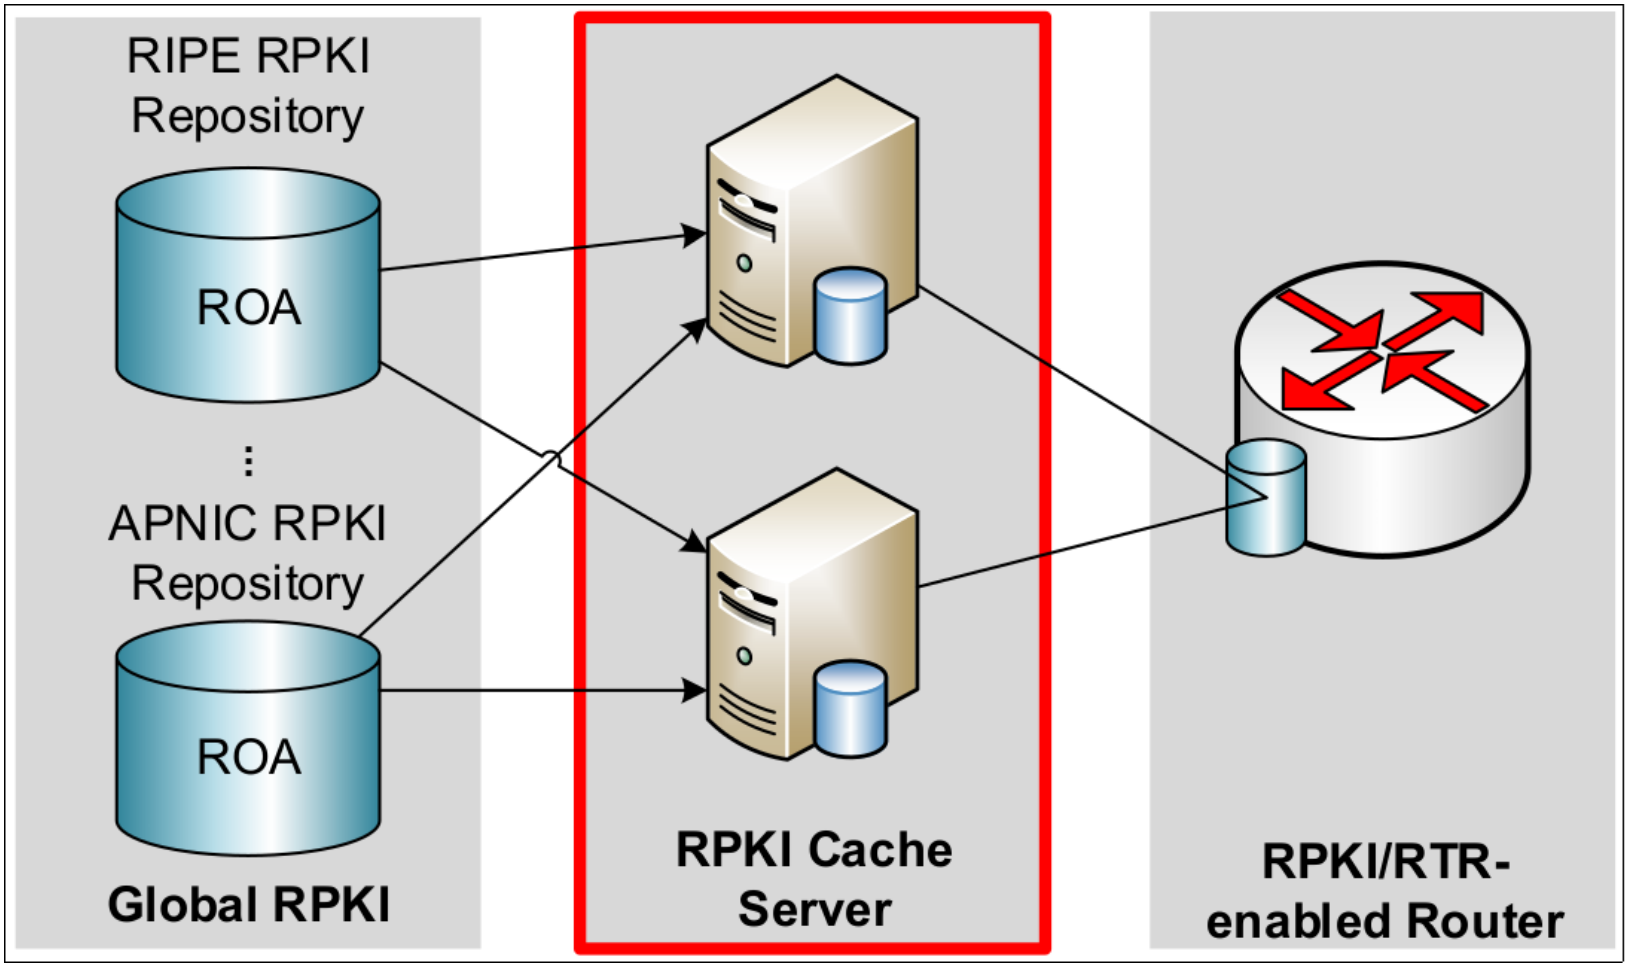

Research Question

- Are RPKI cache servers consistent?
  - Why or why not?
  - How does this effect validation?

Methodology

1. Collect trust anchor repository access logs
2. Synthesize access method and server identities
3. Analyze

Synchronization Access Method

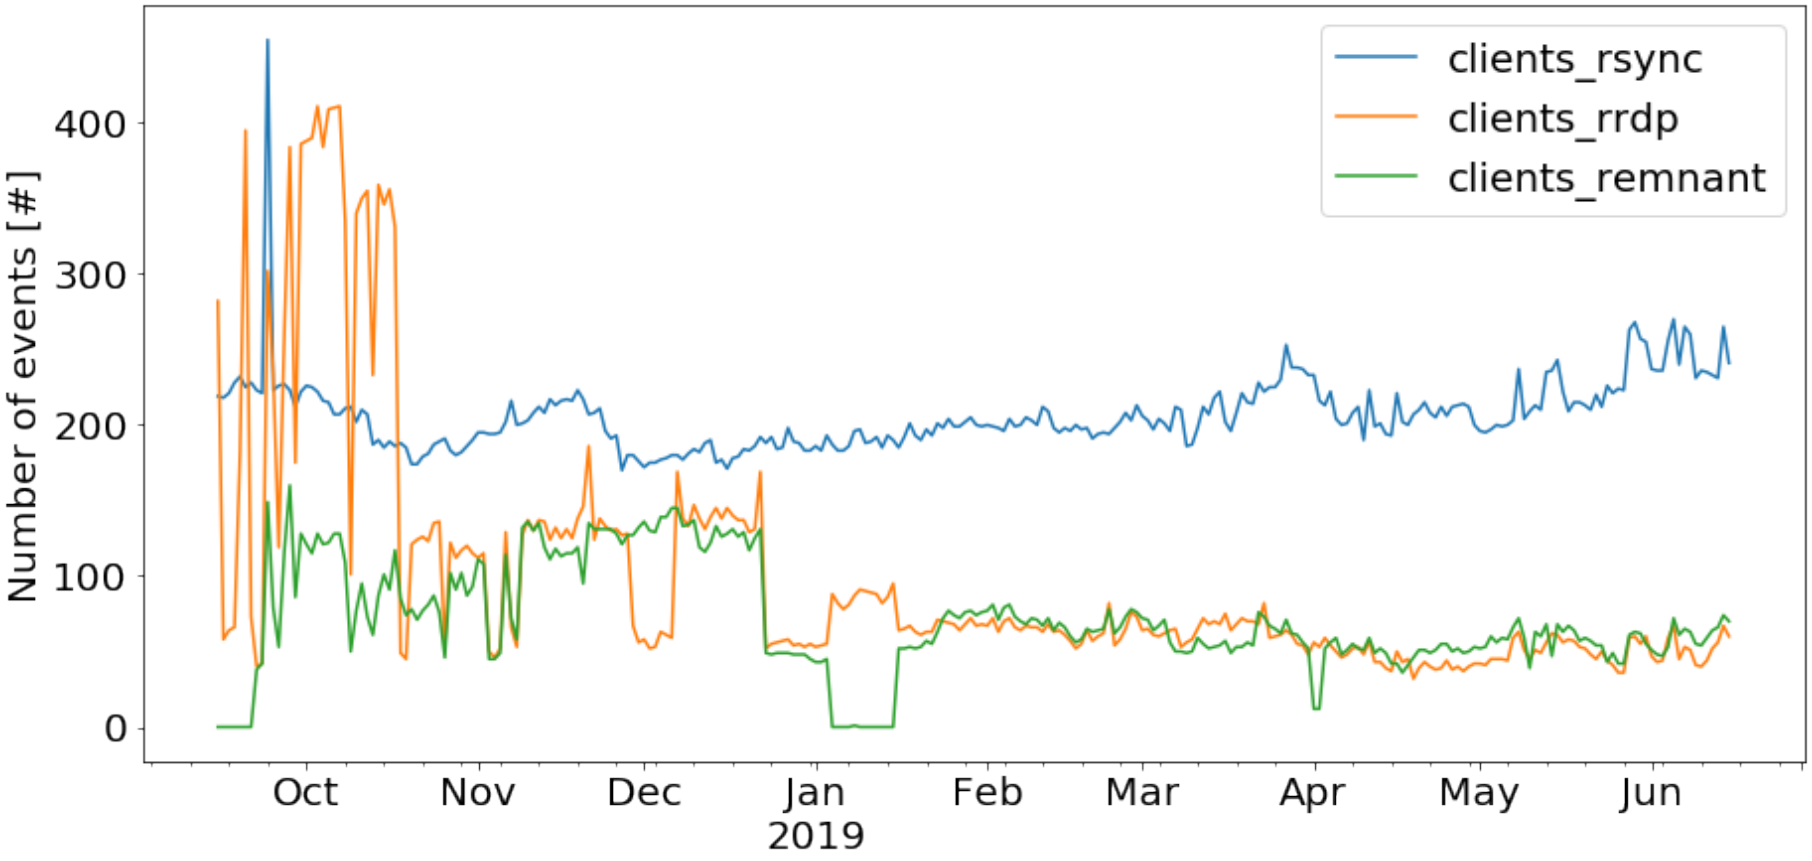

RIR Accesses

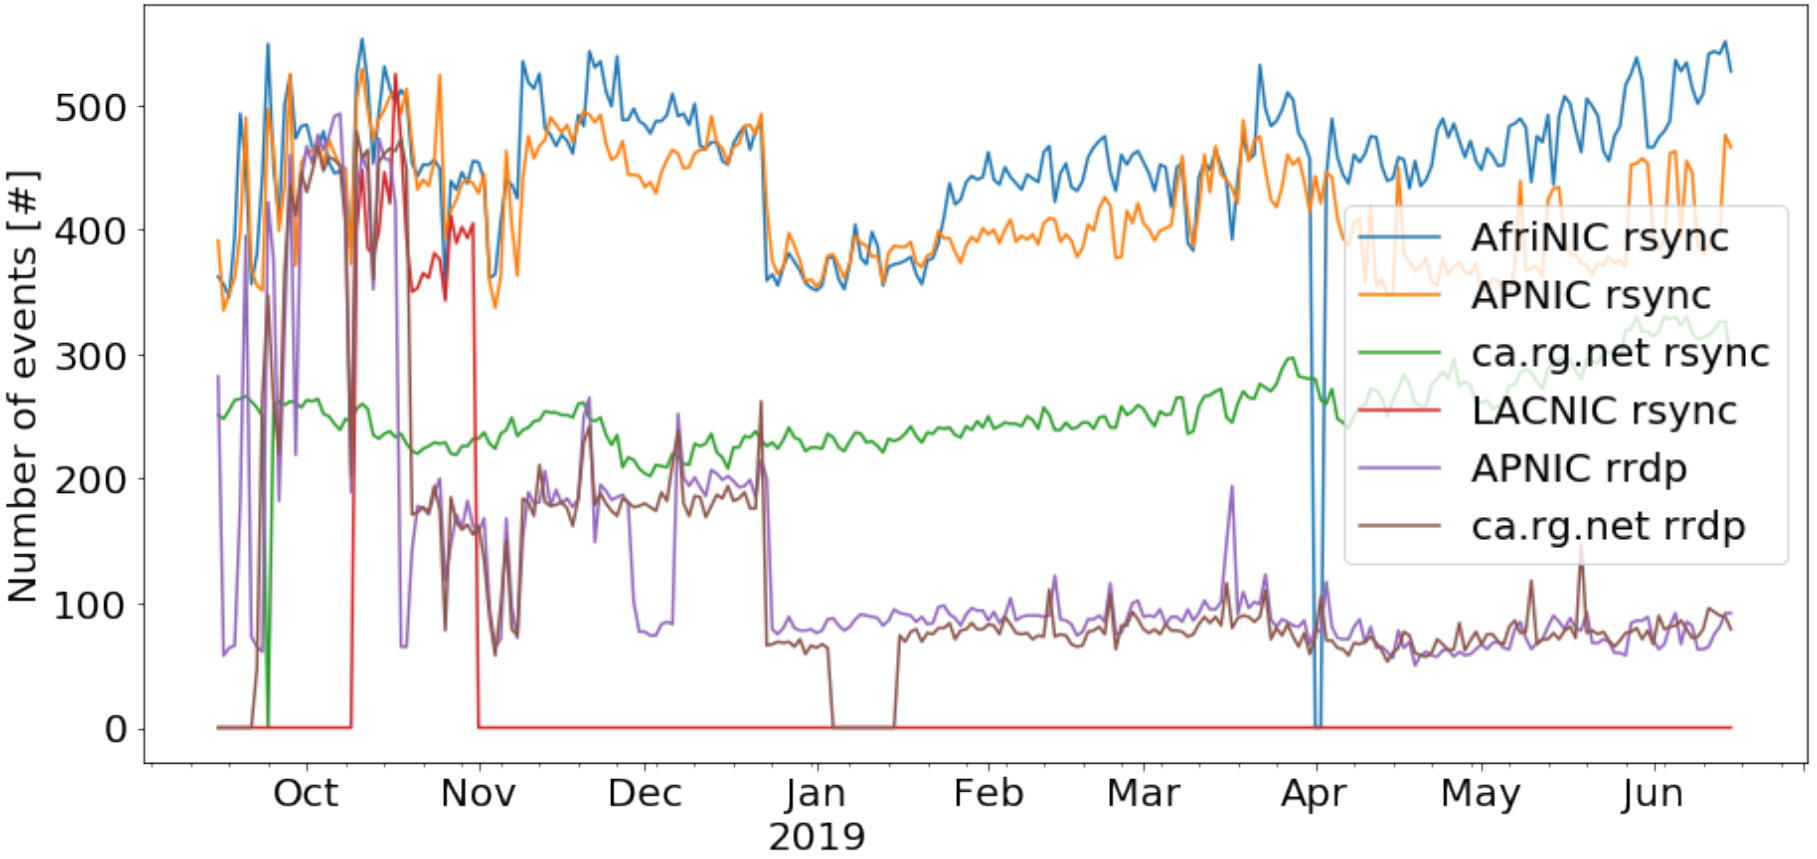

RPKI Cache Server Population

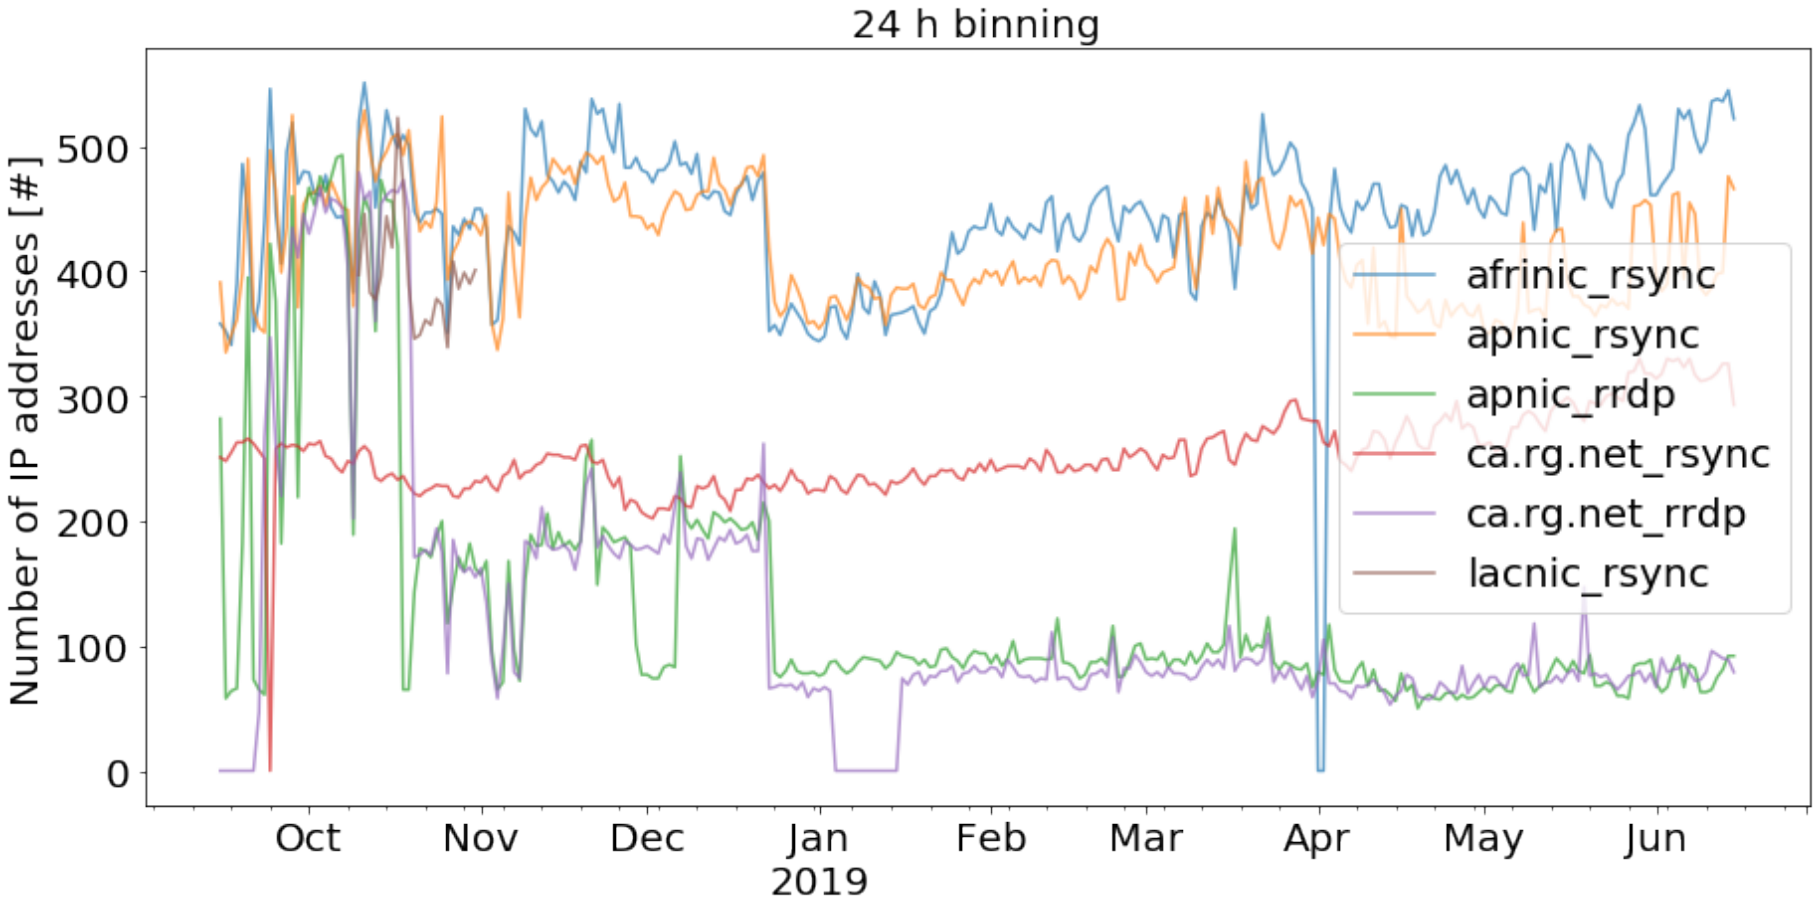

Cache Server Differences

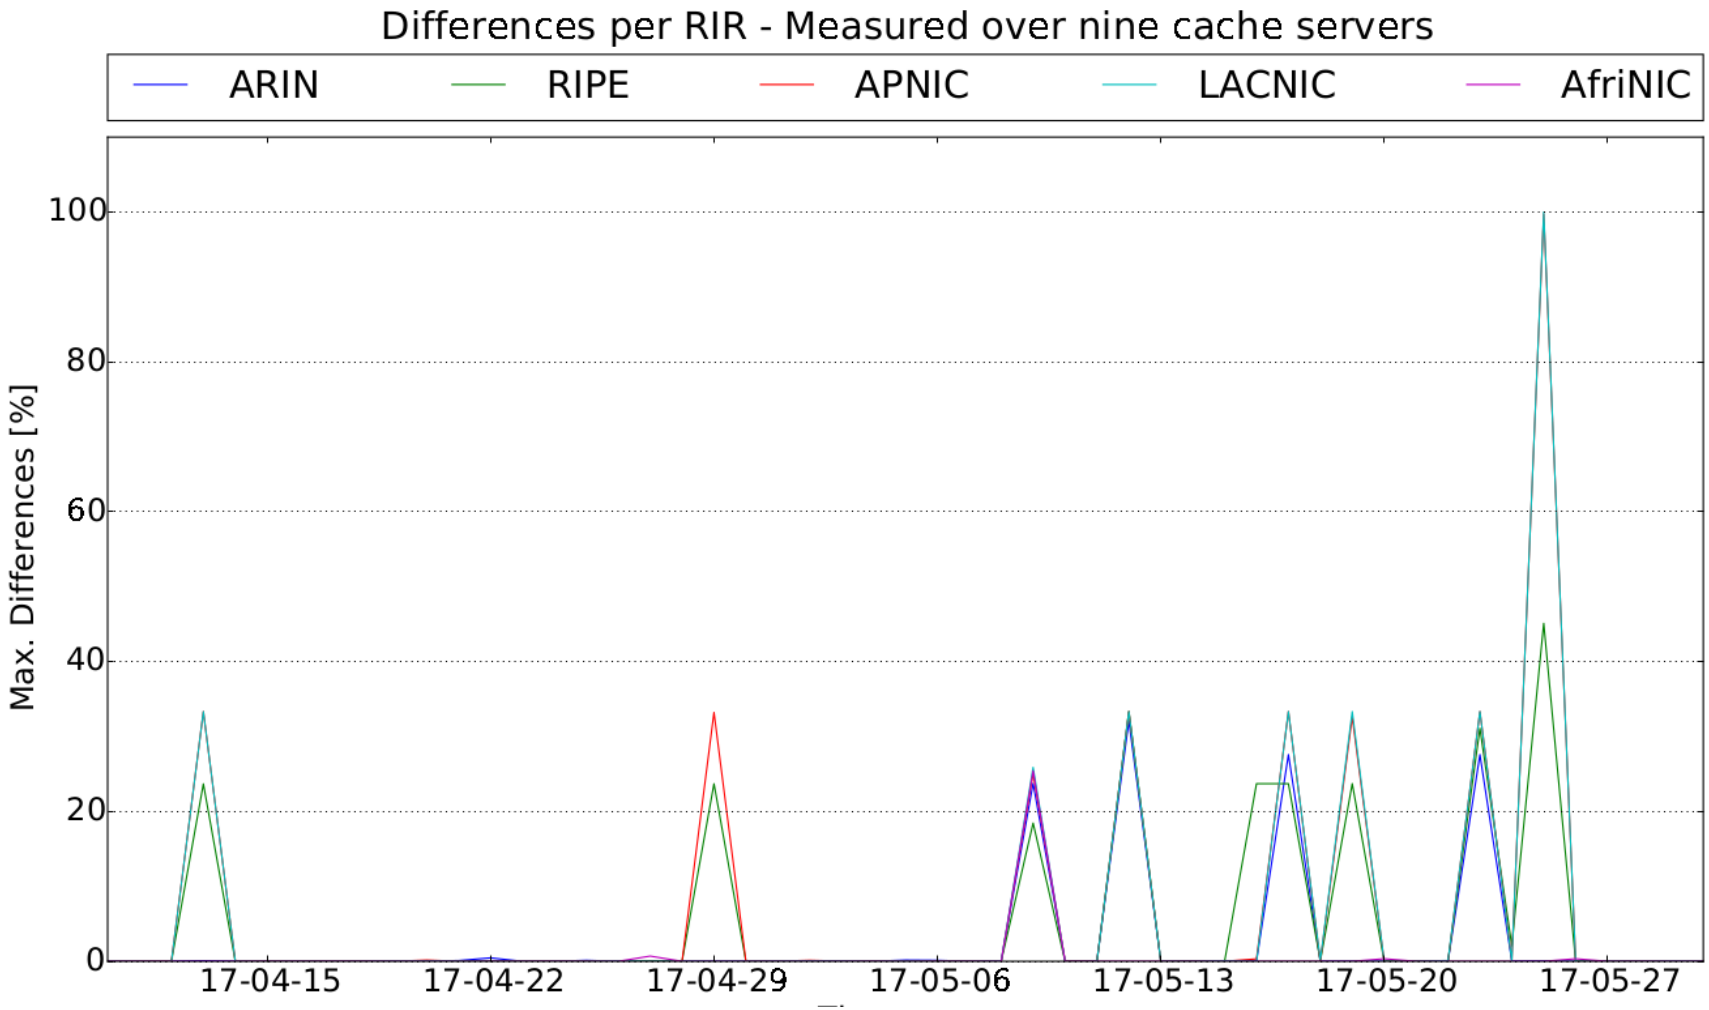

Auxiliary Work: RPKI Cache and Router Discovery

- Internet Survey
  - rpki-rtr / TCP / 323
  - rpki-rtr TLS / TCP / 324
  - Routinator / TCP / 8282 (documentation example)
  - Junos / TCP / 2222 / “no problem here"
- Passive DNS survey
  - rpki* , [.-]rpki*
- rsync traffic flows to trust anchors

Contact Info

- John Kristoff
- Email: jtk@depaul.edu
- WWW: https://aharp.iorc.depaul.edu
- GitHub: https://github.com/jtkristoff
- Twitter: https://twitter.com/jtkristoff