


# Load-Balancing as a Network Primitive


Nikhil Handigol, Srini Seetharaman, Mario Flajslik, Aaron Gember,  
 Vjekoslav Brajkovic, Ramesh Johari, Nick McKeown, Guru Parulkar, Aditya  
 Akella, Arvind Krishnamurthy, Nick Feamster  
 Stanford University, DT Labs, U. Wisconsin, U. Washington, Georgia Tech.



## Goal: Load-balancing web requests in wide-area networks

### Load-balancing in the wide-area

- The problem
  - Which server? Which path?
- Important problem
  - Critical for all scale-out services
  - Big \$\$
- Hard problem
  - Scale -1000s of servers, millions of clients, high requests/sec
  - Multiple ingress points

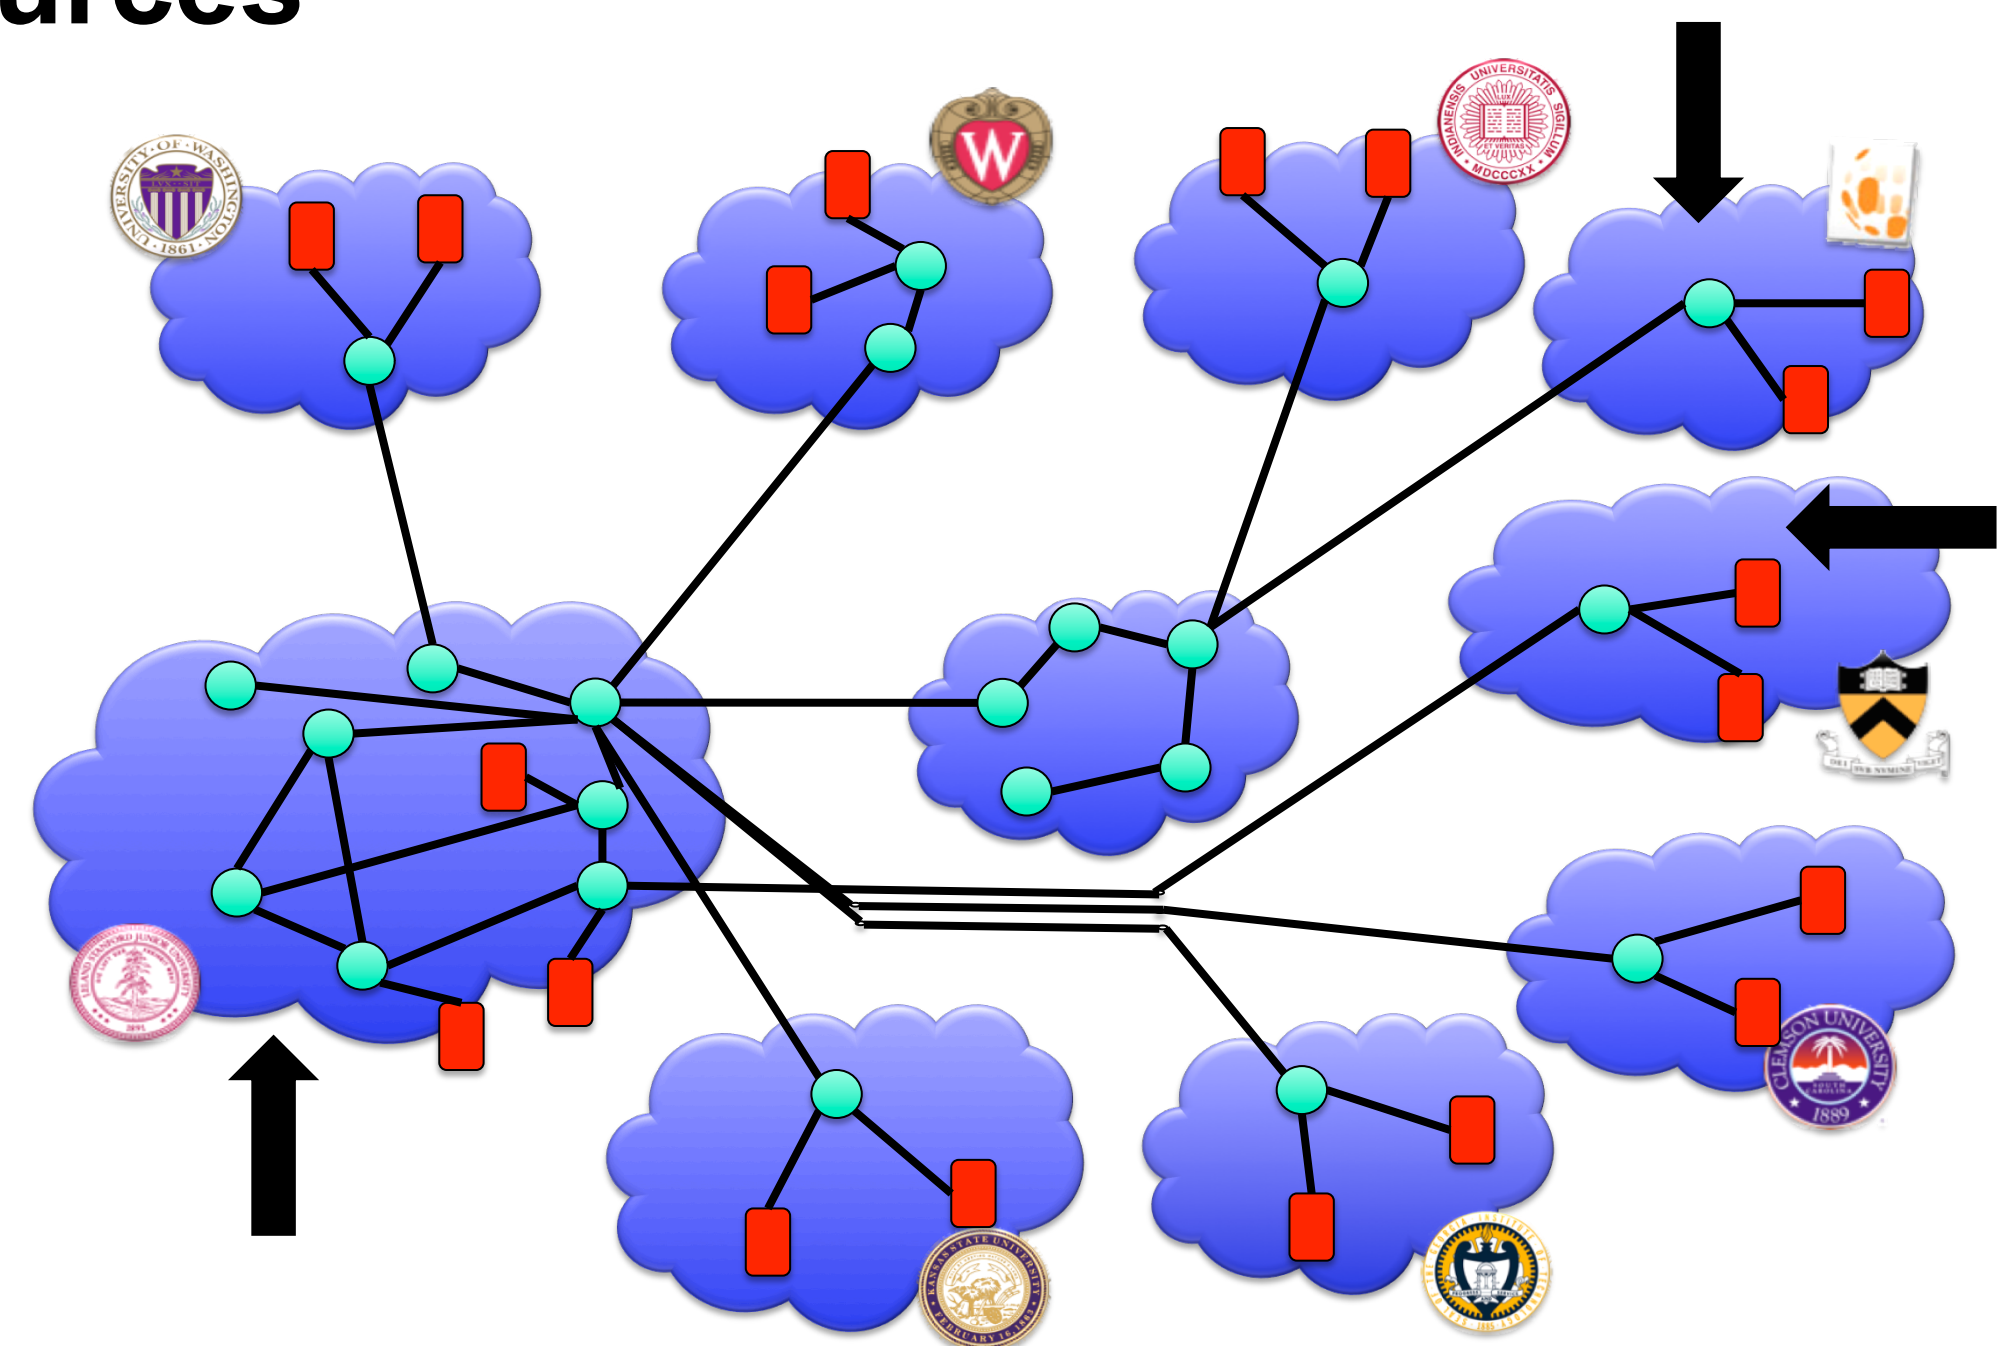


### Research Objective

- Today:
  - Server selection – CDN
  - Path selection – Network
- Ideal scheme:
  - Joint (server, path) selection
- How much worse are today's disjoint (server, path) selection schemes compared to the "ideal" scheme?


### GENI Resources

- PlanetLab-based computation substrate interconnected by OpenFlow network spanning 9 campuses
  - Stanford University
  - University of Washington
  - University of Wisconsin
  - Indiana University
  - GPO Lab
  - Princeton University
  - Clemson University
  - Georgia Tech.
  - Kansas State University

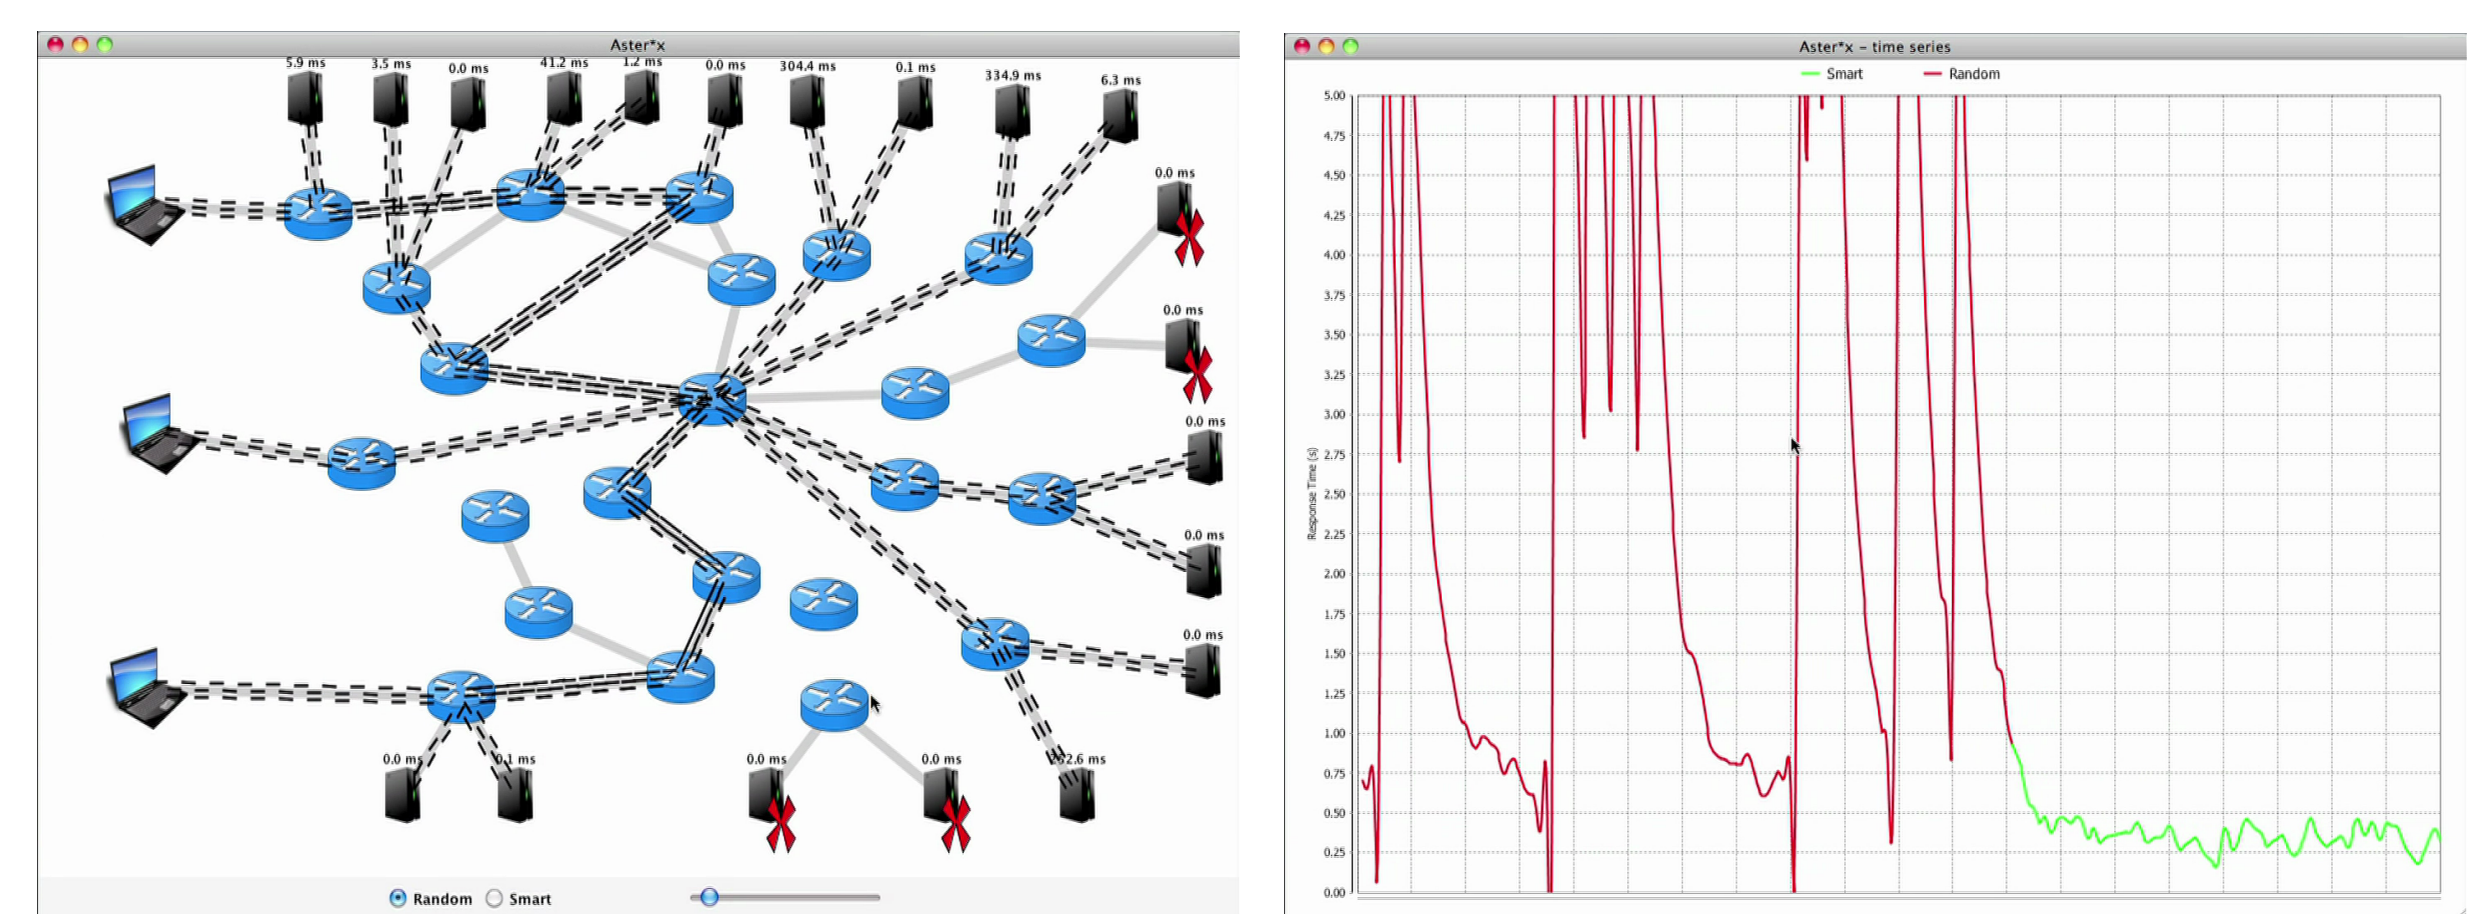


### Experiments and Results

- Topologies – BRITE, CAIDA, Rocketfuel
- Mininet-RT emulation
- Load-balancing schemes
  1. Greedy server selection + shortest-path
  2. Greedy server selection + TE
  3. Joint (server, path) selection
- Results
  - Scheme 1: At least 2x worse over 50% times
  - Scheme 2: Close to *ideal*!



### GENI Experiments



- Future: Extensive evaluation on the GENI slice

### Demos and Publications

- Demos: GEC6, GEC8, GEC9
- Paper: under submission.