Cooperative Content Delivery

Ingmar Poese, Nikhil Handigol, Kang Chen

Motivation

- 3 parties involved in content distribution
  - CDN
  - ISP
  - Consumers
- Only the CDN does server selection

- Each party wants to optimize its own objective

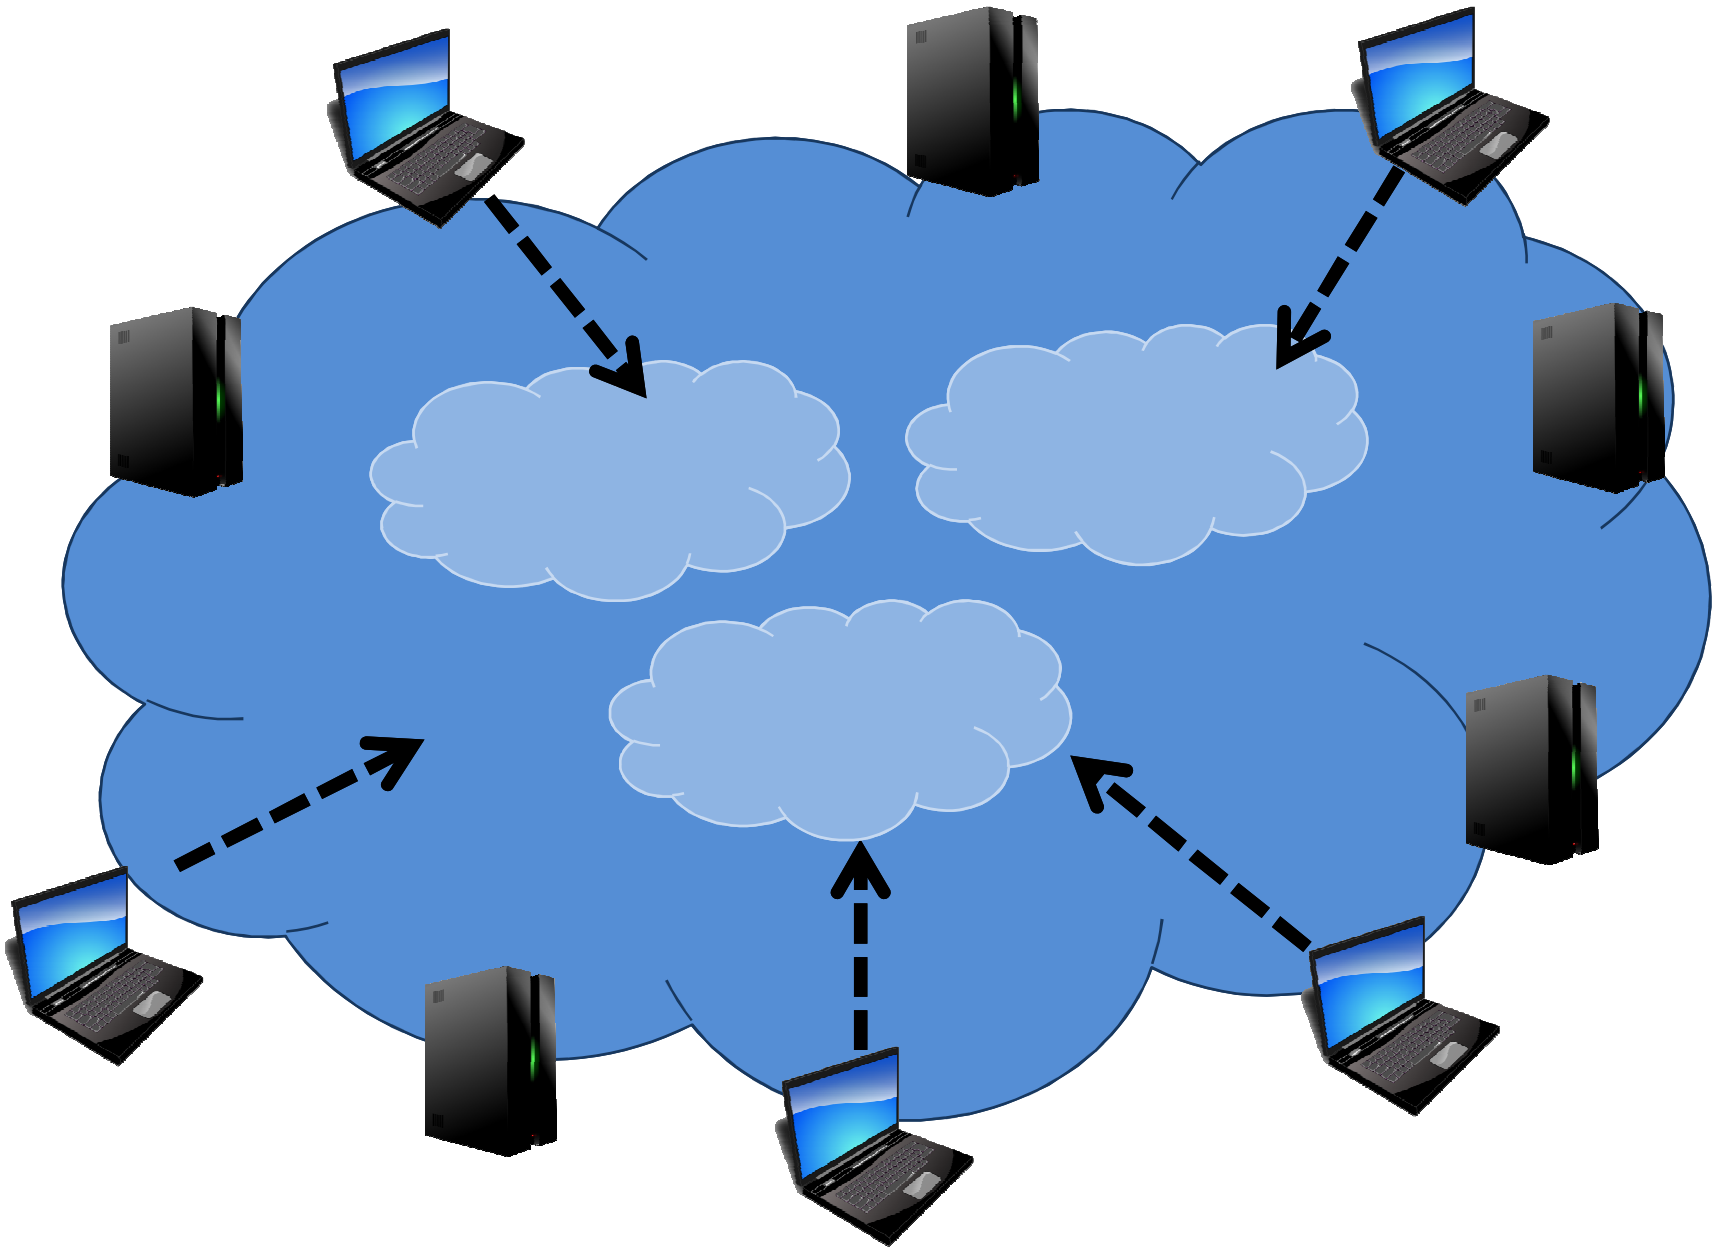

Relevance

- CDNs generate a large fraction of Internet traffic
- Host massive scale-out services
  - Google, Yahoo!, Facebook
- Server selection
  - Need for an evaluation of CDN performance
  - Need for new schemes to improve performance content delivery

Research Problem

- Can cooperation benefit the parties involved in content distribution?
- What information needs to be shared?
- How to share this information?

Experiments

- Spawn an ISP-like network with servers and clients
- GENI/GLAB Resources
  - Core Network – custom topology setup
  - Servers/consumers – PlanetLab, ProtoGENI, etc.
- Evaluation of 4 CDN server-selection cases
  - CDN
  - ISP
  - Consumer
  - Collaborative scheme

Expected Results

- Collaborative server selection
  - Pareto optimal
    - no party can get better performance without hurting others
  - Improved performance of all involved parties

International Collaboration

- Access to resources spanning GENI and GLAB
- Joint publications
- Potential collaboration on other research topics

Preliminary Ideas

- CDN Metrics
  - Server load
  - Application characteristics
  - Content popularity
- Network Metrics
  - Load (delay, utilization, jitter)
  - Location
  - Topology information
  - Connectivity information to other networks
- User Metrics
  - Application in use
  - (Quality of experience)

Protocol design

- Users communicate their application/content before server selection happens
- CDN-ISP and ISP-ISP negotiation to aid server selection and content migration

The Problem Today

- CDNs do server selection, with incomplete knowledge
  - Network state hidden
  - User content not specified
- 3 entities involved in the process
  - CDN
  - ISP
  - Client