McCormick
Northwestern Engineering
Department of Electrical Engineering and Computer Science

Greening the Internet with Nano Data Centers

V. Valancius, N. Laoutaris, L. Massoulié, C. Diot and P. Rodriguez

Presented by: Mario Sanchez
Instructor: Fabian Bustamante
Date: 02/21/2011

Data Center Limitations

- Over-Provisioning
  - Match peak demand
  - Redundancy requirements
- High-cost heat dissipation
  - Heat dissipation accounts for at least 20% to 50% of total power consumption
- Increased distance to end-users
  - Increases bandwidth-mileage requirements
  - Adds to the energy consumption of networking equipment
- Expensive to build/deploy

Data Centers Today

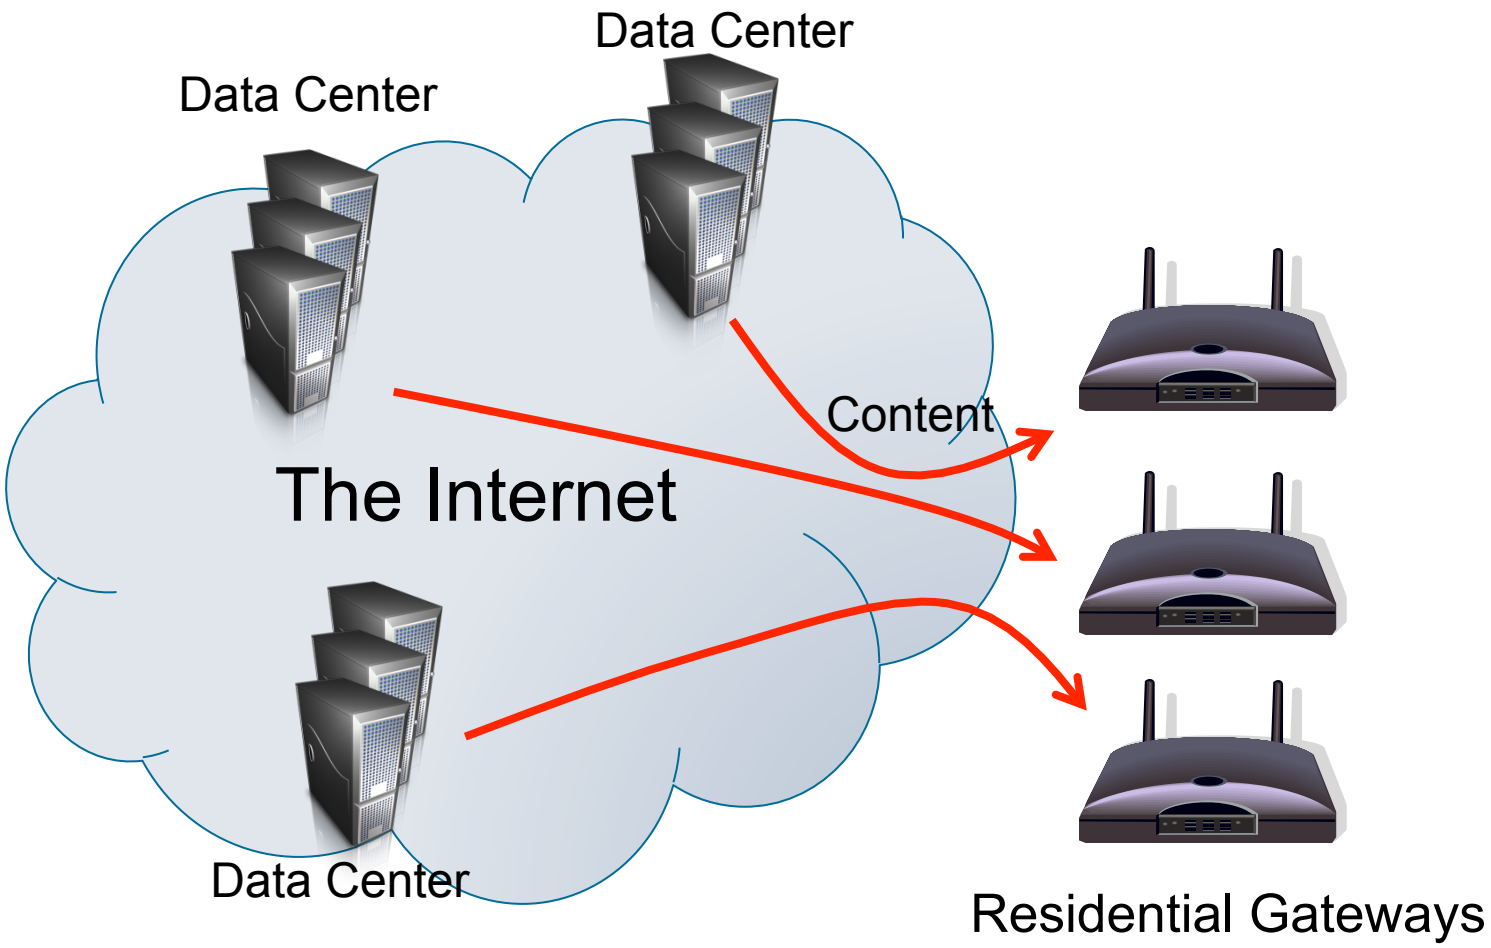

Nano Data Center Model (NaDa)

- Distributed service platform based on tiny managed "servers" located at the edges of the network
- Follows P2P philosophy, with one main difference:

coordinated & **managed** by a single entity (ISP)

- Contributions:
  - Develop model to evaluate energy consumption of Data Centers vs NaDa
  - Apply model in the context of video-on-demand services
  - Use trace-driven simulation to quantify energy savings

Nano Data Center Model (NaDa)

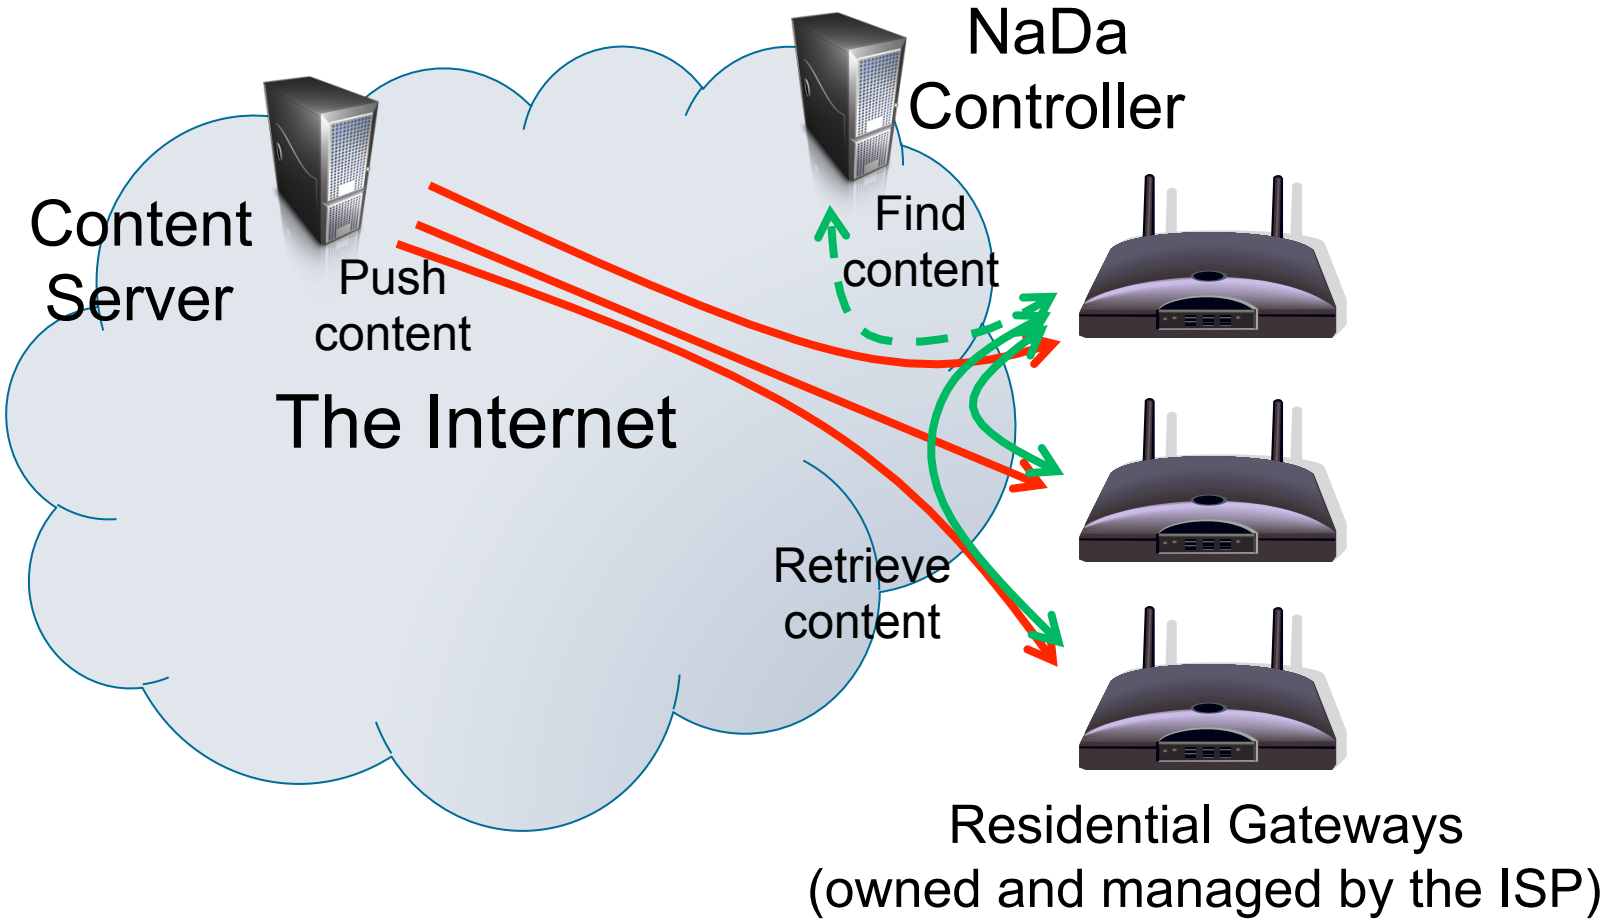

Advantages - NaDa

- Heat dissipation
  - **Power Usage Efficiency (PUE):** ratio between total power consumed by a data center and the power actually delivered to it IT equipment
- Service proximity
  - Information travels shorter distances
  - Reduced energy in powering networking equipment that carries traffic
- Self-scalability
  - Growth organically with the network
- Energy efficiency
  - No baseline powering is paid

Energy Use in Servers

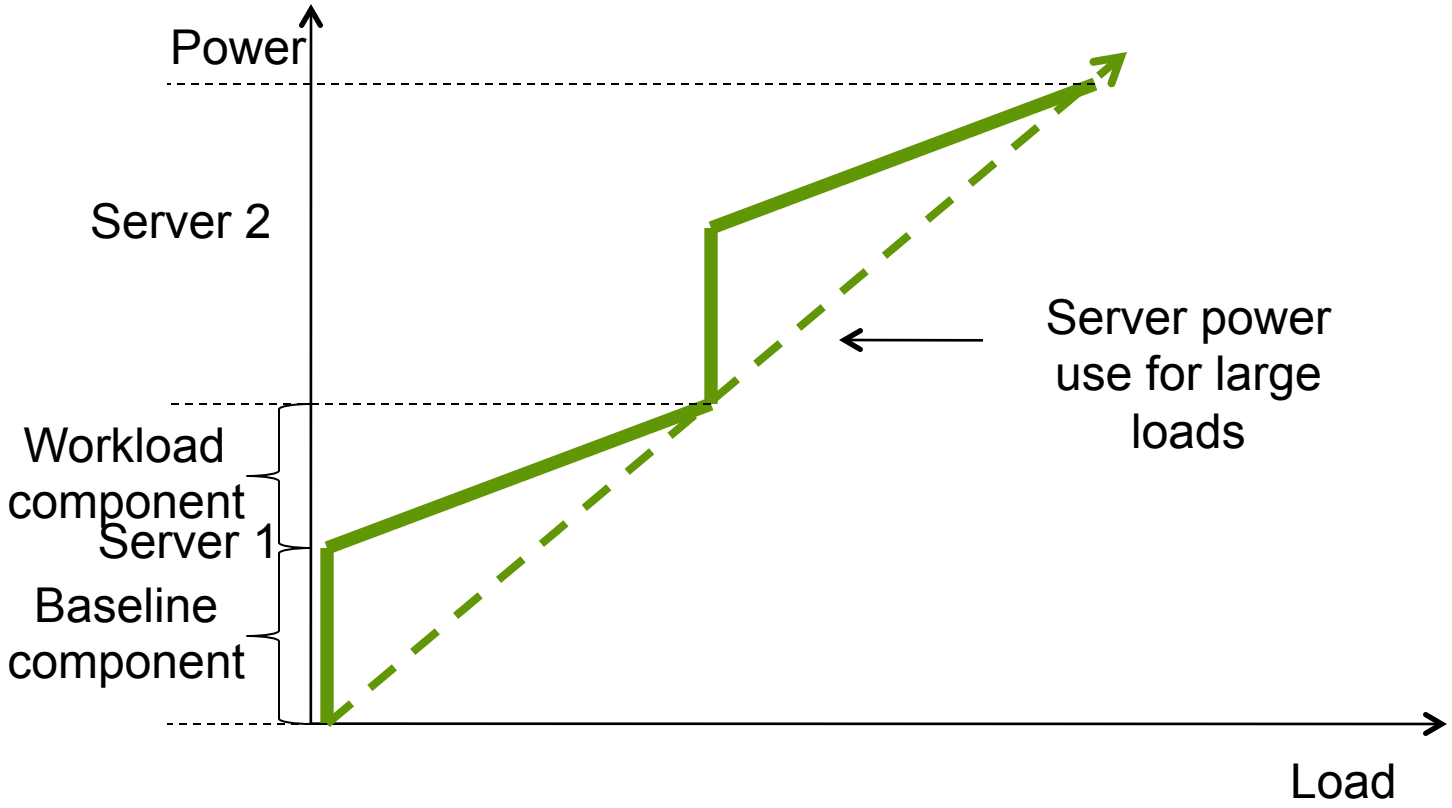

DSL Gateways, discount baseline

- If we dedicate home gateways for NaDa service
  - 50-80% of power is consumed when idle (baseline power)
  - **Less** efficient than server

- If we use gateways only when they are active
  - Power consumption is proportional to load
  - **More** efficient than dedicated server!

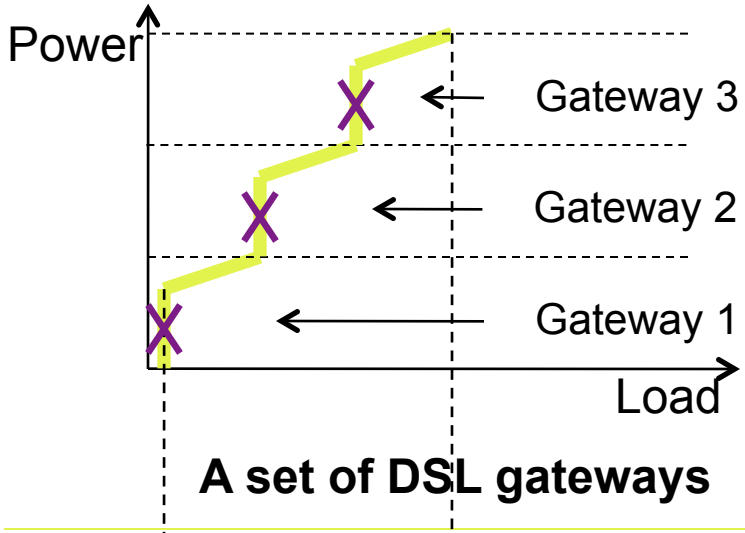

A set of DSL gateways

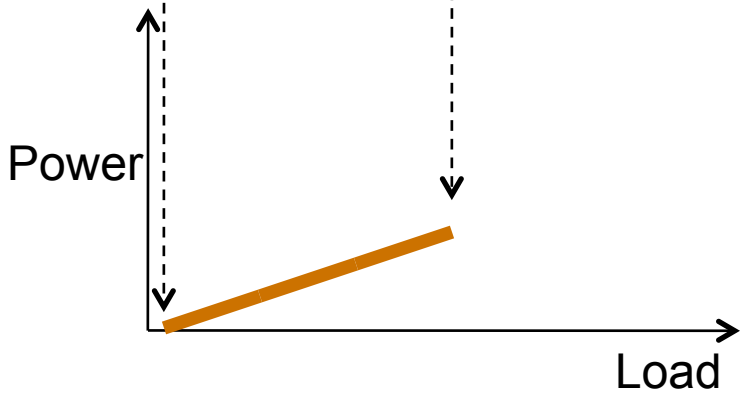

DSL gateways, discounting baseline

Energy Savings Model

Power consumption = b + ax

- b: base line consumption
- a: slope of load-dependent consumption
- x: traffic (Mbps)

- Carrier grade IPTV streaming server
  - Server energy use ~211Joules/Gbit
- Thomson residential DSL gateways
  - Gateway energy use: ~100Joules/Gbit
- Cisco router chasis
  - Network energy use: ~150Joules/Gbit

Estimated Energy use

- Data center power use estimation:
  - Data center PUE = 1.2
- NaDa home gateway power use estimation
  - Residential GW PUE = 1.07

- Upper bound on energy savings:

$$100\% - \frac{100 \cdot 1.07 \cdot x}{211 \cdot 1.2 \cdot x} \approx 55\%$$

Traces/Simulation - Dataset

a) **Netflix movie database:** number of rental, 90 minutes duration
b) **Youtube traces:** view count and length
c) **IPTV access statistics:** random sample 2k users, program duration

VoD & Simulation

- Three key components
  - Gateways: provide storage, bandwidth resources
  - The tracker: coordinates all VoD activities in NaDa
  - Content servers: provide content to gateways an clients (if needed)

- Large metropolitan area:
  - Each user assumed to have identical network distance to every other user in the network
  - Every user has the same distance to the content servers

Content placement strategy

- Split content in smaller chunks
- Replicate content across random gateways
- Optimization problem: determine replication number for each movie
- **Hot-warm-cold placement method**

  (Movies are partitioned into three groups):
  - Popular movies **=>** replicated on all gateways.
  - Subsequent most popular movies => replicated minimally
  - Less popular movies => not stored within NaDa

Traces/Simulation

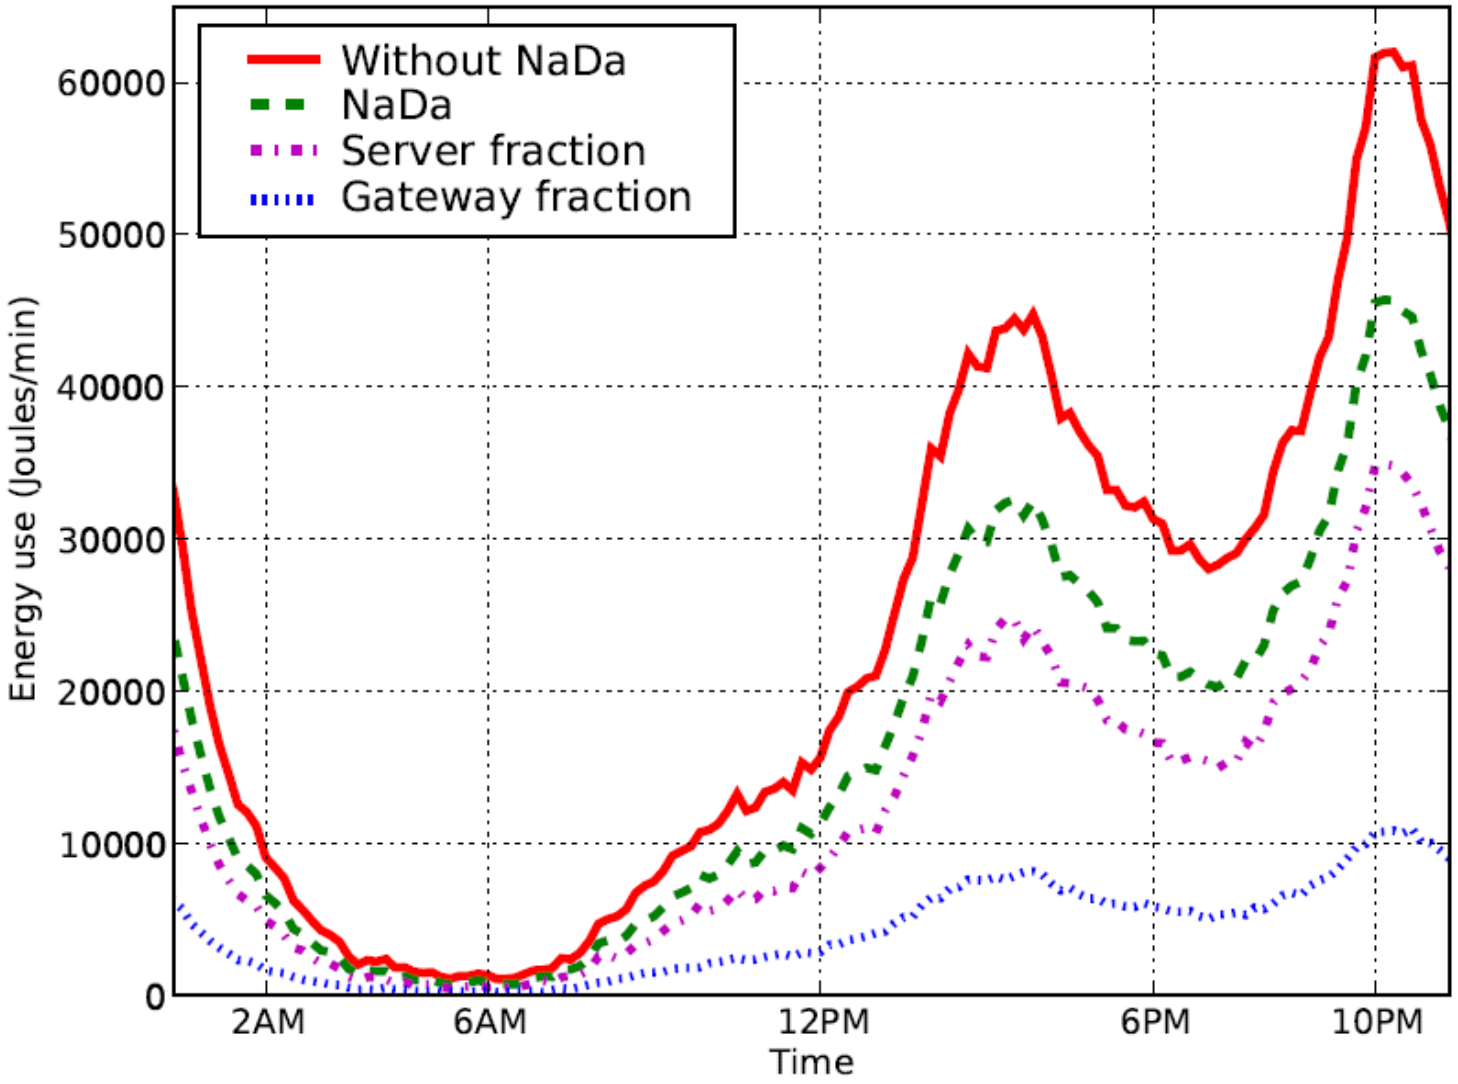

Effect of Storage

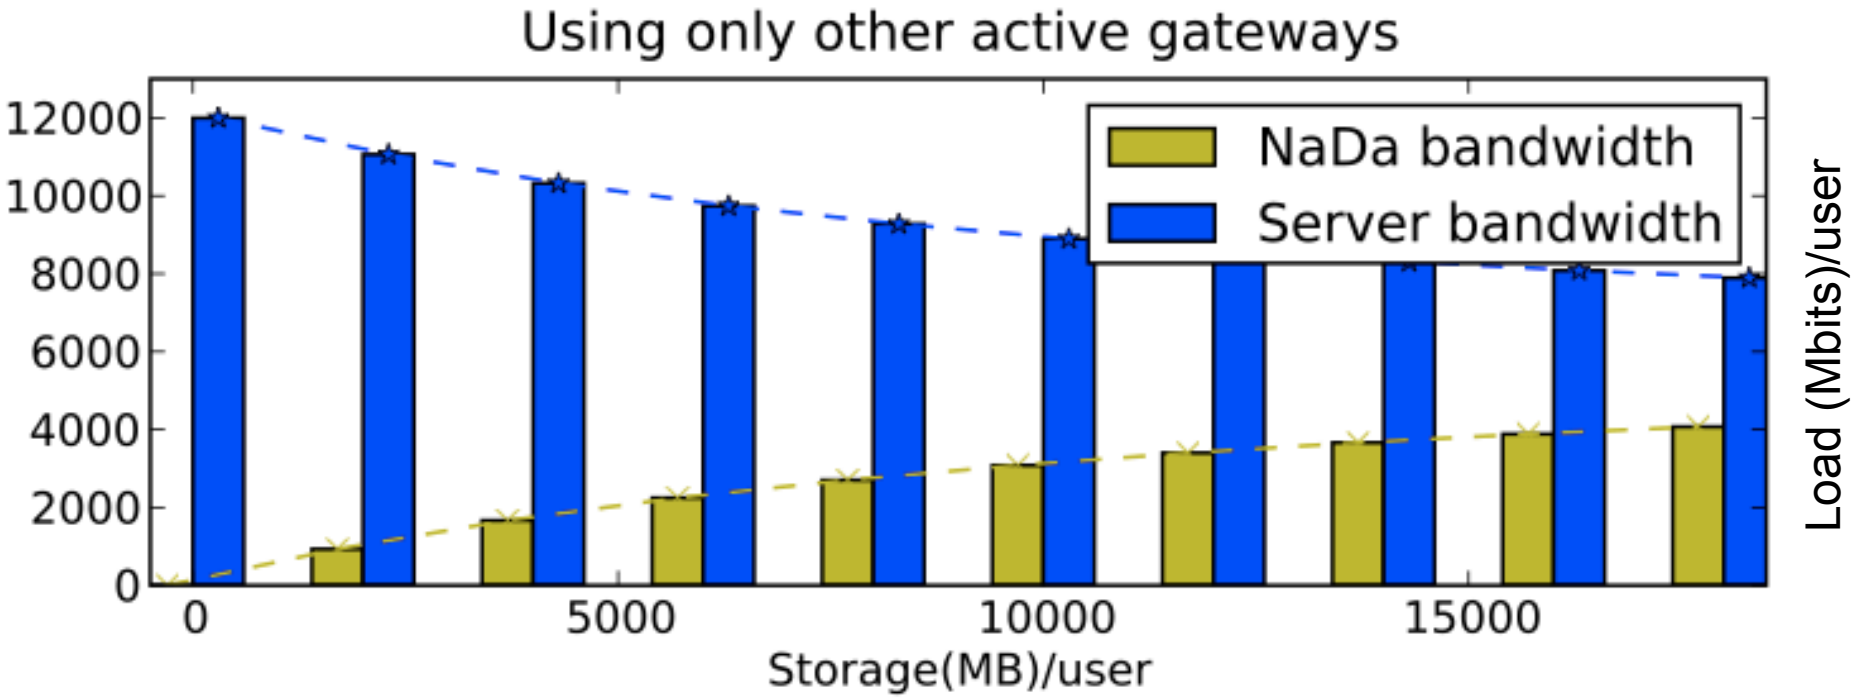

Effect of Video Rates

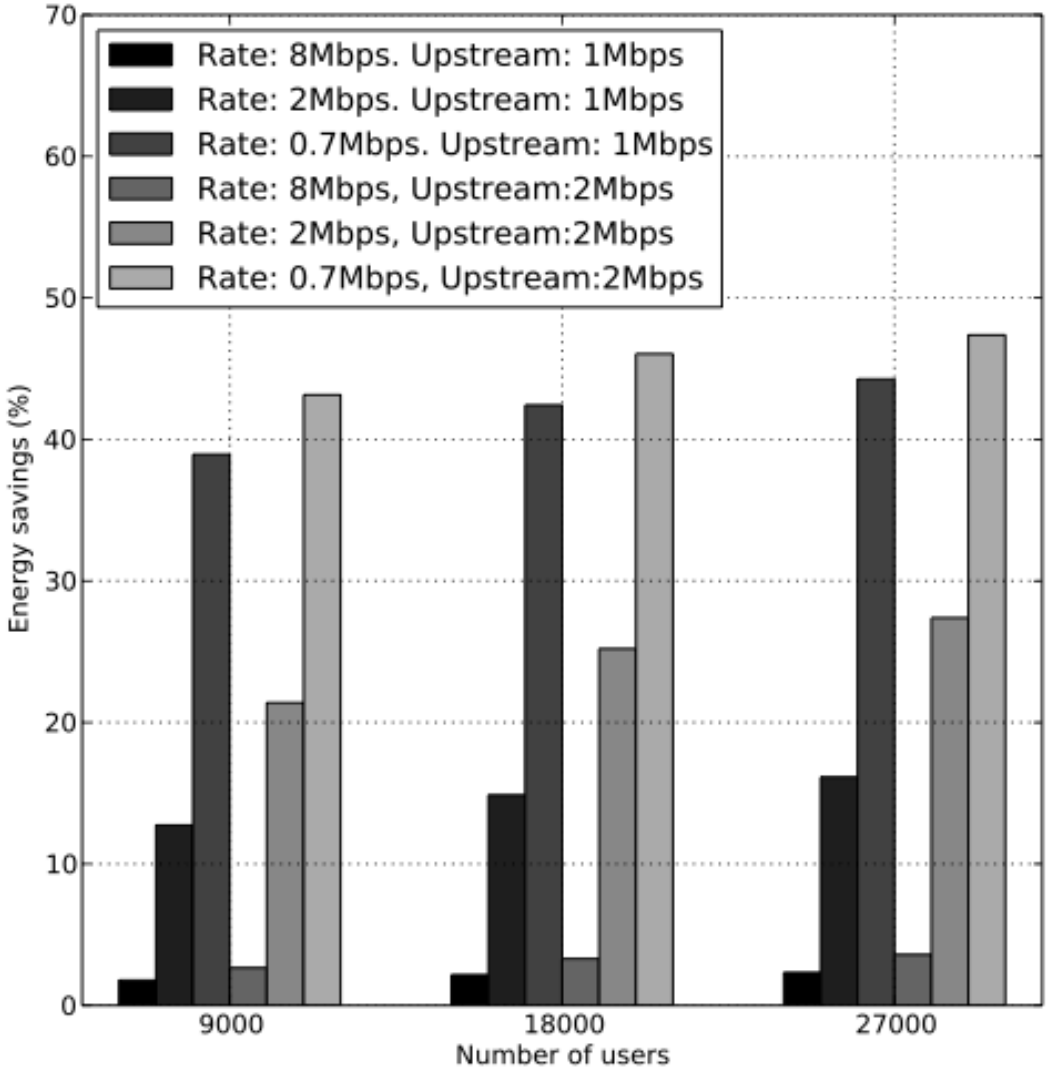

Effect of Network

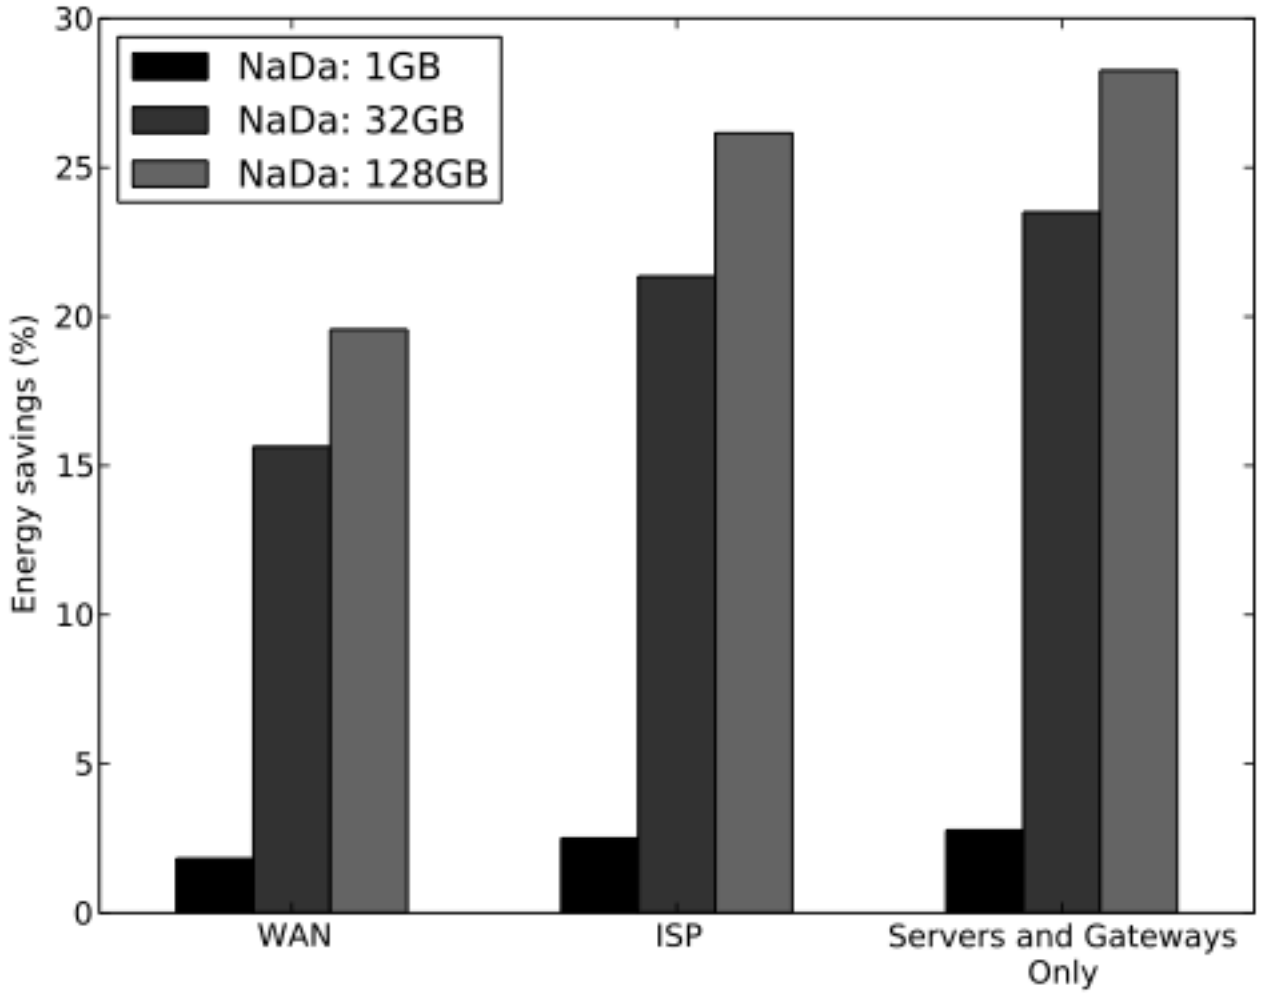

Discussion