Content-peering Dynamics of Autonomous Caches in a Content-centric Network

Valentino Pacifici, György Dán

Laboratory for Communication Networks
School of Electrical Engineering
KTH, Royal Institute of Technology
Stockholm - Sweden

Stockholm, December 13, 2012

Content-centric Networks

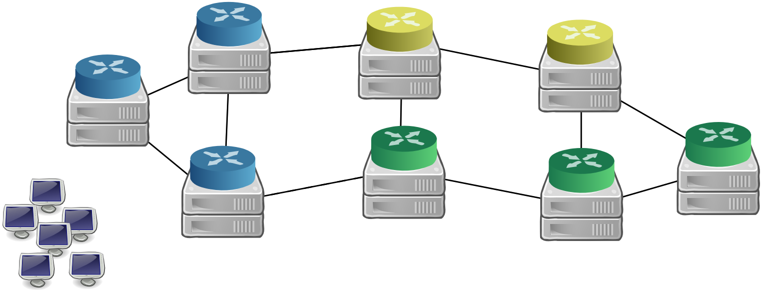

- Caches part of the protocol stack
- Existing research optimizes *global performance*
  - Cache dimensioning
  - Efficient routing
  - Efficient cache eviction policies

Content-centric Networks

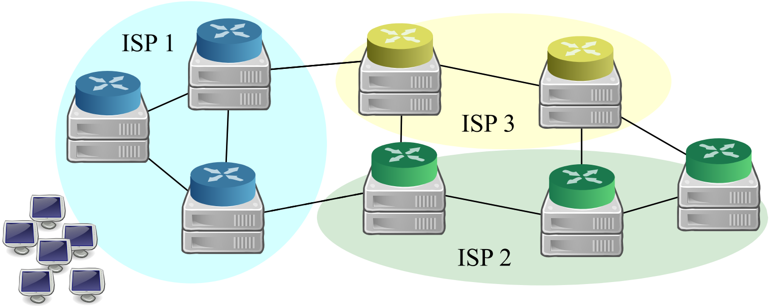

- Caches part of the protocol stack
- Existing research optimizes *global performance*
  - Cache dimensioning
  - Efficient routing
  - Efficient cache eviction policies

Content-centric Networks

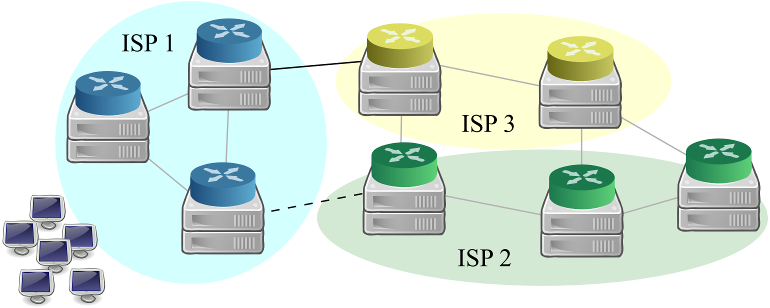

- Caches part of the protocol stack
- Existing research optimizes *global performance*
  - Cache dimensioning
  - Efficient routing
  - Efficient cache eviction policies

Content-centric Networks

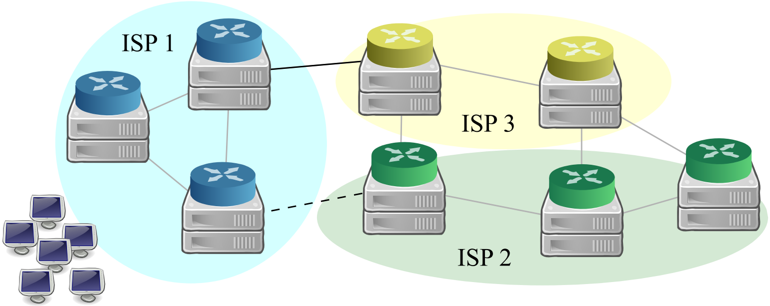

- Networks of caches optimized for *local performance*
- Decrease transit traffic costs through *content-level peering*
- New challenges:
  - Stability of cache content
  - Coordination among ASes
  - Effect of eviction

Modeling the Interaction among ASes

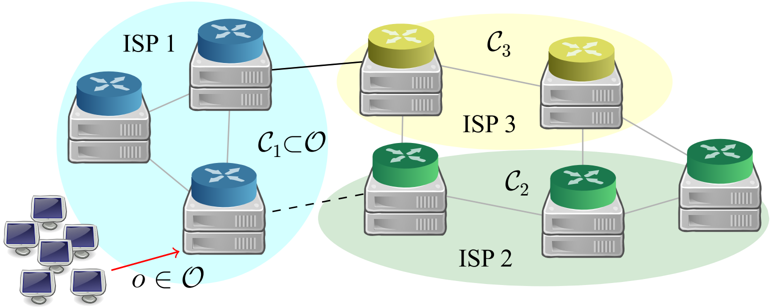

Modeling the Interaction among ASes

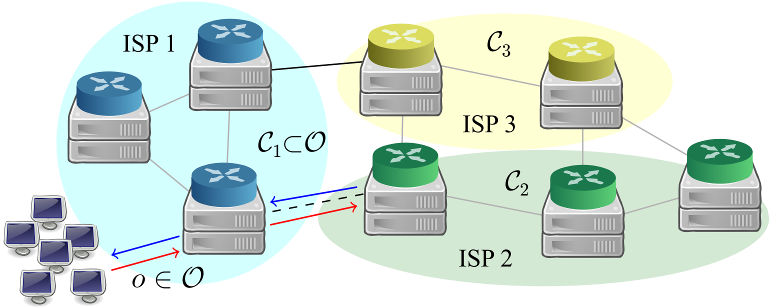

Modeling the Interaction among ASes

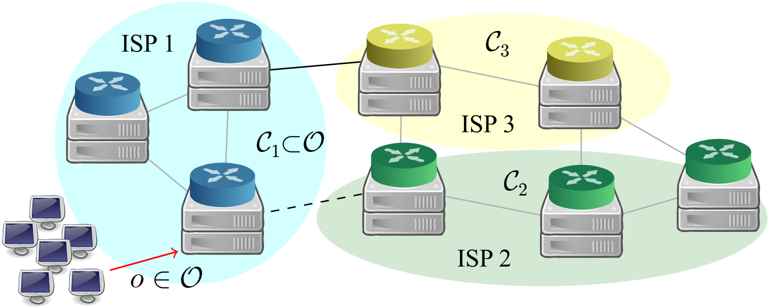

Modeling the Interaction among ASes

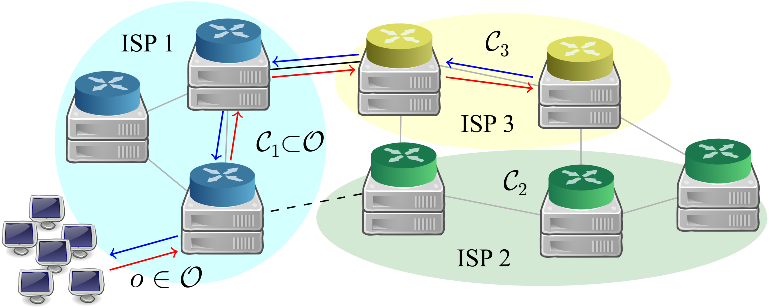

Modeling the Interaction among ASes

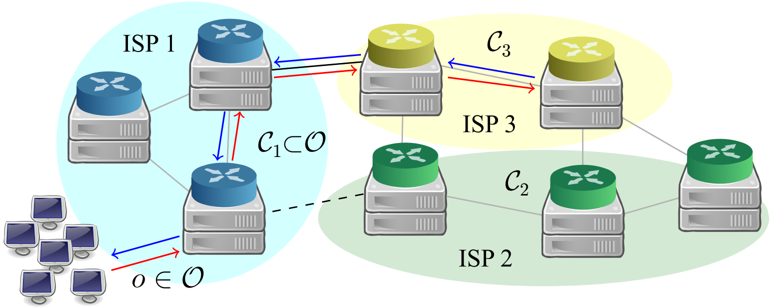

- Each ISP optimizes its internal network through
    - Routing of content and interest messages
    - Cache dimensioning and eviction policies

Modeling the Interaction among ASes

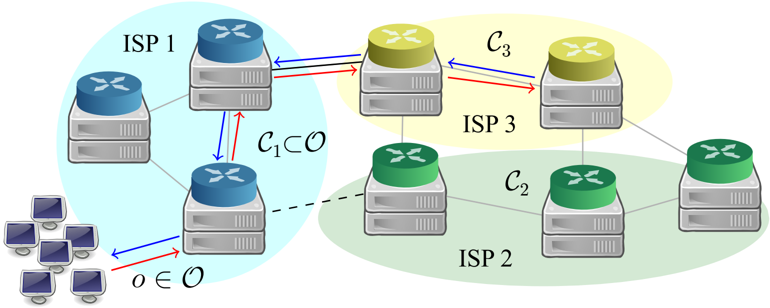

- Each ISP optimizes its internal network through
  - Routing of content and interest messages
  - Cache dimensioning and eviction policies

$$\Downarrow$$

- $\mathcal{L}_i = \mathcal{H}_i \cup \mathcal{C}_i \rightarrow$ content available at ISP i
- $\mathcal{R}_i = \bigcup_{j \in \mathcal{N}(i)} \mathcal{L}_j \rightarrow$ content available from the peering ISPs

Modeling the Interaction among ASes

- Interest messages (i.m.) for item o at ISP i
  - $w_i^o \in \mathbb{R}_+ \rightarrow$ average arrival intensity
  - Independent Reference Model (IRM):
    - inter arrival times $\sim F_i^o(x) = 1 - e^{-w_i^o x}$

Modeling the Interaction among ASes

- Interest messages (i.m.) for item o at ISP i
  - $w_i^o \in \mathbb{R}_+ \rightarrow$ average arrival intensity
  - Independent Reference Model (IRM):
    - inter arrival times $\sim F_i^o(x) = 1 - e^{-w_i^o x}$
- Unit cost of serving item o at ISP i
  - α_i if o available locally or at peering ISP ($o \in \mathcal{L}_i \cup \mathcal{R}_i$)
  - γ_i if o retrieved from a transit link ($o \notin \mathcal{L}_i \cup \mathcal{R}_i$)

Modeling the Interaction among ASes

- Interest messages (i.m.) for item o at ISP i
  - $w_i^o \in \mathbb{R}_+ \rightarrow$ average arrival intensity
  - Independent Reference Model (IRM):
    - inter arrival times $\sim F_i^o(x) = 1 - e^{-w_i^o x}$
- Unit cost of serving item o at ISP i
  - α_i if o available locally or at peering ISP ($o \in \mathcal{L}_i \cup \mathcal{R}_i$)
  - γ_i if o retrieved from a transit link ($o \notin \mathcal{L}_i \cup \mathcal{R}_i$)

Total cost for ISP i:

$$C_i(\mathcal{C}_i, \mathcal{C}_{-i}) = \alpha_i \sum_{\mathcal{L}_i \cup \mathcal{R}_i} w_i^o + \gamma_i \sum_{\mathcal{O} \setminus \{\mathcal{L}_i \cup \mathcal{R}_i\}} w_i^o,$$

Coordinated Content-peering

- Peering ISPs periodically exchange *Summary cache* $\mathcal{L}_i$
- Upon interest message for item o at ISP i:
  - if $o \in \mathcal{L}_i \rightarrow$ the item is served
  - if $o \in \mathcal{R}_i \rightarrow$ i.m. forwarded to a peer
  - if $o \notin \mathcal{R}_i \rightarrow$ i.m. forwarded to transit provider

Coordinated Content-peering

- Peering ISPs periodically exchange *Summary cache* $\mathcal{L}_i$
- Upon interest message for item o at ISP i:
  - if $o \in \mathcal{L}_i \rightarrow$ the item is served
  - if $o \in \mathcal{R}_i \rightarrow$ i.m. forwarded to a peer
  - if $o \notin \mathcal{R}_i \rightarrow$ i.m. forwarded to transit provider

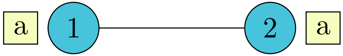

Figure: No content-peering

Coordinated Content-peering

- Peering ISPs periodically exchange *Summary cache* $\mathcal{L}_i$
- Upon interest message for item o at ISP i:
  - if $o \in \mathcal{L}_i \rightarrow$ the item is served
  - if $o \in \mathcal{R}_i \rightarrow$ i.m. forwarded to a peer
  - if $o \notin \mathcal{R}_i \rightarrow$ i.m. forwarded to transit provider

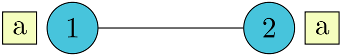

Figure: No content-peering

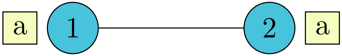

Figure: Content-peering

Coordinated Content-peering

- Peering ISPs periodically exchange *Summary cache* $\mathcal{L}_i$
- Upon interest message for item o at ISP i:
  - if $o \in \mathcal{L}_i \rightarrow$ the item is served
  - if $o \in \mathcal{R}_i \rightarrow$ i.m. forwarded to a peer
  - if $o \notin \mathcal{R}_i \rightarrow$ i.m. forwarded to transit provider

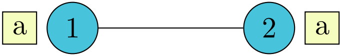

Figure: No content-peering

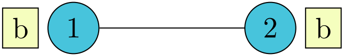

Figure: Content-peering

Coordinated Content-peering

- Peering ISPs periodically exchange *Summary cache* $\mathcal{L}_i$
- Upon interest message for item o at ISP i:
  - if $o \in \mathcal{L}_i \rightarrow$ the item is served
  - if $o \in \mathcal{R}_i \rightarrow$ i.m. forwarded to a peer
  - if $o \notin \mathcal{R}_i \rightarrow$ i.m. forwarded to transit provider

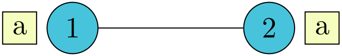

Figure: No content-peering

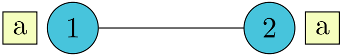

Figure: Content-peering

Coordinated Content-peering

- Peering ISPs periodically exchange *Summary cache* $\mathcal{L}_i$
- Upon interest message for item o at ISP i:
  - if $o \in \mathcal{L}_i \rightarrow$ the item is served
  - if $o \in \mathcal{R}_i \rightarrow$ i.m. forwarded to a peer
  - if $o \notin \mathcal{R}_i \rightarrow$ i.m. forwarded to transit provider

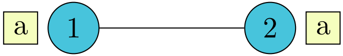

Figure: No content-peering

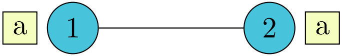

Figure: Content-peering

- Need for algorithms to reach stable allocation

Coordinated Content-peering

- Peering ISPs periodically exchange *Summary cache* $\mathcal{L}_i$
- Upon interest message for item o at ISP i:
  - if $o \in \mathcal{L}_i \rightarrow$ the item is served
  - if $o \in \mathcal{R}_i \rightarrow$ i.m. forwarded to a peer
  - if $o \notin \mathcal{R}_i \rightarrow$ i.m. forwarded to transit provider

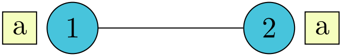

Figure: No content-peering

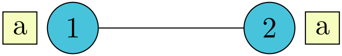

Figure: Content-peering

- Need for algorithms to reach stable allocation

PERFECT INFORMATION: Perfect estimation of w_i^o

Cache-or-Wait (CoW) Algorithm

- Independent set $\mathcal{I} \subseteq N$: it does not contain peering ISPs

Cache-or-Wait (CoW) Algorithm

- Independent set $\mathcal{I} \subseteq N$: it does not contain peering ISPs

At every time slot t:

- Pick $\mathcal{I}_t$.
- ISPs $i \in \mathcal{I}_t$ allowed to change their cached items $\mathcal{C}_i(t-1) \to \mathcal{C}_i(t)$
- For all $j \notin \mathcal{I}_t$, $\mathcal{C}_j(t) = \mathcal{C}_j(t-1)$.
- Inform peering ISPs about $\mathcal{C}_i(t)$

- Transitions as a Markov Process P^0

Cache-or-Wait (CoW) Algorithm

- Independent set $\mathcal{I} \subseteq N$: it does not contain peering ISPs

At every time slot t:

- Pick $\mathcal{I}_t$.
- ISPs $i \in \mathcal{I}_t$ allowed to change their cached items $\mathcal{C}_i(t-1) \to \mathcal{C}_i(t)$
- For all $j \notin \mathcal{I}_t$, $\mathcal{C}_j(t) = \mathcal{C}_j(t-1)$.
- Inform peering ISPs about $\mathcal{C}_i(t)$

- Transitions as a Markov Process P^0

CoW *terminates in an equilibrium allocation after a finite number of* ***efficient updates***

- At slot t no ISP $j \notin \mathcal{I}_t$ can perform an update

Cache-no-Wait (CnW) Algorithm

At every time slot t

- $\forall$ ISP $i \in N$ allowed to perform efficient updates $\mathcal{C}_i(t-1) \to \mathcal{C}_i(t)$
- $\forall$ ISP i informs the ISPs $j \in \mathcal{N}(i)$ about $\mathcal{C}_i(t)$.

Cache-no-Wait (CnW) Algorithm

At every time slot t

- $\forall$ ISP $i \in N$ allowed to perform efficient updates $\mathcal{C}_i(t-1) \rightarrow \mathcal{C}_i(t)$
- $\forall$ ISP i informs the ISPs $j \in \mathcal{N}(i)$ about $\mathcal{C}_i(t)$.

CnW *terminates in an equilibrium allocation with probability one.*

At every time slot:

- Prob. $e^{-w_i^o \Delta}$ item o not requested
- Prob. $\epsilon(\mathcal{C}_i(t-1)) > 0$ no "useful" item requested
- Prob. $k(\mathcal{C}(t-1))$ that updating ISPs belong to an independent set

$$k(\mathcal{C}(t-1)) > \prod_{i \in \mathcal{I}_t} [1 - \epsilon_i(\mathcal{C}_i(t-1))] \cdot \prod_{i \in N \setminus \mathcal{I}_t} \epsilon_i(\mathcal{C}_i(t-1)) > 0$$

Numerical Results - CoW vs CnW

- CAIDA graph: largest connected component in CAIDA dataset
  - 616 ISPs, average node degree 9.66
- Erdős-Rényi (ER) and Barabási-Albert (BA) random graphs
- Arrival intensities $\sim$ Zipf's law with exponent 1
- Cache size 10 at every ISP

Numerical Results - CoW vs CnW

- CAIDA graph: largest connected component in CAIDA dataset
    - 616 ISPs, average node degree 9.66
- Erdős-Rényi (ER) and Barabási-Albert (BA) random graphs
- Arrival intensities ∼ Zipf's law with exponent 1
- Cache size 10 at every ISP

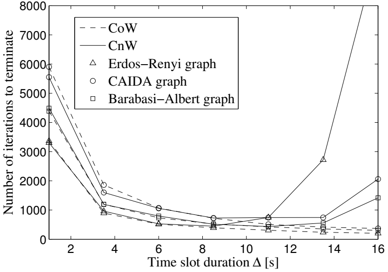

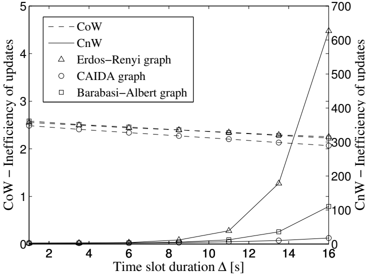

Content-peering under Imperfect Information

- Avg. arrival intensities w_i^o are estimated
- Probability of misestimation:

$$\text{If } w_i^o > w_i^p \;\Rightarrow\; P(\overline{w}_i^o < \overline{w}_i^p) \propto \epsilon e^{-\frac{1}{\beta}(w_i^o - w_i^p)}$$

- System leaves equilibrium allocations
- P^β regular perturbed Markov process of P^0

Content-peering under Imperfect Information

- Avg. arrival intensities w_i^o are estimated
- Probability of misestimation:

$$\text{If } w_i^o > w_i^p \ \Rightarrow \ P(\overline{w}_i^o < \overline{w}_i^p) \propto \epsilon e^{-\frac{1}{\beta}(w_i^o - w_i^p)}$$

- System leaves equilibrium allocations
- P^β regular perturbed Markov process of P^0

When $\beta \to 0$, the stationary distribution of P^β is a stationary distribution of P^0 $\Rightarrow$ equilibrium allocation under perfect information.

- Some cache allocations more likely than others..

Imperfect Information - Disjoint Interests

- Disjoint interests case:

the K_i items with highest arrival intensity of the ISPs form disjoint sets.

The stationary distribution of P^{β} is the stable allocation C^{*} in which every ISP caches its most popular items.

Imperfect Information - Disjoint Interests

- Disjoint interests case:

the K_i items with highest arrival intensity of the ISPs form disjoint sets.

The stationary distribution of P^β is the stable allocation C^* in which every ISP caches its most popular items.

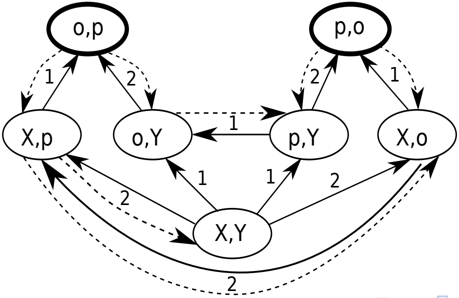

Imperfect Information - Numerical Results

- 10^5 time slots
- 50 ISPs, $K_i = 1$ for every $i \in N$
- τ estimation interval in seconds for a LFU cache

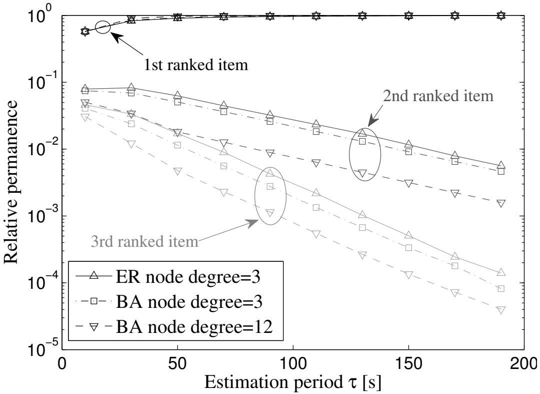

Conclusions

- Model of interaction between ISPs in CCNs
- Content-level peering reaches a stable allocation
- Fast convergence if no simultaneous updates by peering ISPs

Conclusions

- Model of interaction between ISPs in CCNs
- Content-level peering reaches a stable allocation
- Fast convergence if no simultaneous updates by peering ISPs

- Incorrect estimate of content popularity
  - No stable state
  - Still cost-efficient cache allocations
  - Insights about most likely cache allocations