Distributed Algorithms for Content Allocation in Interconnected Content Distribution Networks

Valentino Pacifici, György Dán

Laboratory for Communication Networks
KTH Royal Institute of Technology
Stockholm, Sweden

Hong Kong, April 30, 2015

Booming Content Delivery Market

Content Distribution in the Internet

- 2009: P2P traffic → up to 70 % of Internet traffic
- 2013: Netflix + YouTube → 50 % fixed network traffic
- 2017: Video traffic → 80 % of IP traffic

Booming Content Delivery Market

Content Distribution in the Internet

- 2009: P2P traffic → up to 70 % of Internet traffic
- 2013: Netflix + YouTube → 50 % fixed network traffic
- 2017: Video traffic → 80 % of IP traffic

Over-the-top Content providers

- ↑ Quality of delivered content
- ↑ User demand for content
- ↑ REVENUE

Booming Content Delivery Market

Content Distribution in the Internet

- 2009: P2P traffic → up to 70 % of Internet traffic
- 2013: Netflix + YouTube → 50 % fixed network traffic
- 2017: Video traffic → 80 % of IP traffic

Over-the-top Content providers

- ↑ Quality of delivered content
- ↑ User demand for content

- ↑ Revenue

Content Delivery Networks (CDNs)

- ↑ Delivered content on behalf of OTT providers
- ↑ 2017: delivery of $\frac{2}{3}$ of total video traffic

- ↑ Revenue

Network Operator Managed CDNs

Network operators

- Content distribution stresses network infrastructure (OTT, P2P)
- Network operators not part of revenue chain

↓

Network Operator Managed CDNs

Network operators

- Content distribution stresses network infrastructure (OTT, P2P)
- Network operators not part of revenue chain

↓

Network Operator Managed CDNs (nCDNs)

- Storage of content close to the customers
- Objectives:
    1. Improve user's QoE → decrease latency
    2. Decrease traffic cost → decrease network traffic

Network Operator Managed CDNs

Network operators

- Content distribution stresses network infrastructure (OTT, P2P)
- Network operators not part of revenue chain

$\downarrow$

Network Operator Managed CDNs (nCDNs)

- Storage of content close to the customers
- Objectives:
  1. Improve user's QoE $\rightarrow$ decrease latency
  2. Decrease traffic cost $\rightarrow$ decrease network traffic

Content Allocation Problem

- nCDNs periodically update content allocation based on predicted demands
- **Pre-fetching**: nCDN i decides on allocation $A_i \in \mathcal{O}$ and fetches the content

Autonomous Content Delivery Networks

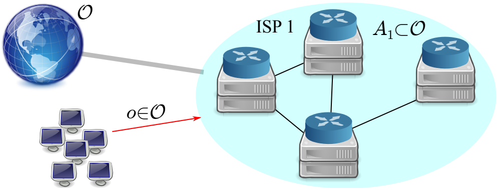

- nCDN optimized for *local performance*

Autonomous Content Delivery Networks

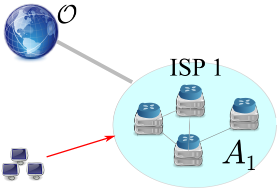

- nCDN optimized for *local performance*

Autonomous Content Delivery Networks

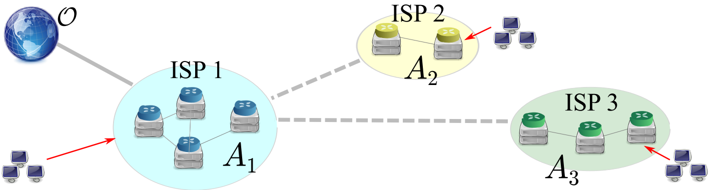

- nCDN optimized for *local performance*

Interconnected nCDNs

- Maximize users' QoE
  - Retrieve items from connected nCDNs with lowest latency

Autonomous Content Delivery Networks

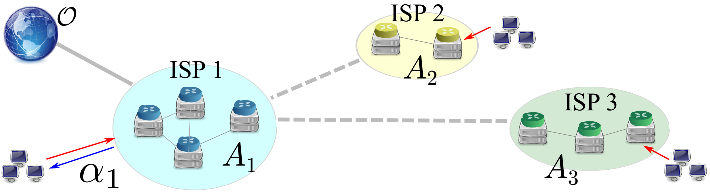

- nCDN optimized for *local performance*

Interconnected nCDNs

- Maximize users' QoE
  - Retrieve items from connected nCDNs with lowest latency

Cost for serving one request (latency)

- α_i from own nCDN i

Autonomous Content Delivery Networks

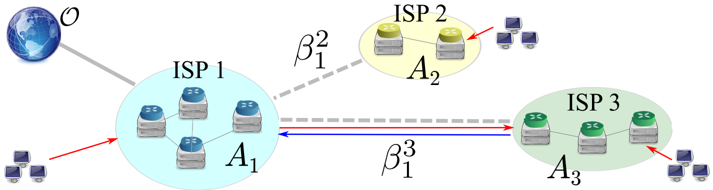

- nCDN optimized for *local performance*

Interconnected nCDNs

- Maximize users' QoE
  - Retrieve items from connected nCDNs with lowest latency

Cost for serving one request (latency)

- α_i from own nCDN i
- β_i^j from connected nCDN j

Autonomous Content Delivery Networks

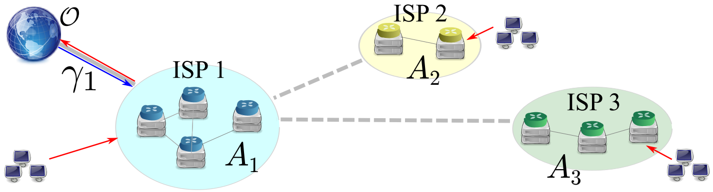

- nCDN optimized for *local performance*

Interconnected nCDNs

- Maximize users' QoE
  - Retrieve items from connected nCDNs with lowest latency

Cost for serving one request (latency)

- α_i from own nCDN i
- β_i^j from connected nCDN j
- γ_i from content provider

$$\boxed{\alpha_i \leq \beta_i^j < \gamma_i}$$

Interconnected nCDNs - The Cost Model

- K_i storage capacity of nCDN i
- $w_i^o \in \mathbb{R}_+$ demand for item $o \in \mathcal{O}$ at nCDN i
- $\mathcal{R}_i = \bigcup_{j \in \mathcal{N}(i)} A_j \rightarrow$ content available from connected nCDNs
- $\beta_i^o(A_{-i}) \triangleq \min_{j \in \mathcal{N}(i)}\{\beta_i^j | o \in A_j\} \rightarrow$ lowest latency to retrieve item o

Interconnected nCDNs - The Cost Model

- K_i storage capacity of nCDN i
- $w_i^o \in \mathbb{R}_+$ demand for item $o \in \mathcal{O}$ at nCDN i
- $\mathcal{R}_i = \bigcup_{j \in \mathcal{N}(i)} A_j \rightarrow$ content available from connected nCDNs
- $\beta_i^o(A_{-i}) \triangleq \min_{j \in \mathcal{N}(i)}\{\beta_i^j | o \in A_j\} \rightarrow$ lowest latency to retrieve item o

Average latency experienced by customers of operator i :

$$C_i(A_i, A_{-i}) = \sum_{A_i} w_i^o \alpha_i$$

Interconnected nCDNs - The Cost Model

- K_i storage capacity of nCDN i
- $w_i^o \in \mathbb{R}_+$ demand for item $o \in \mathcal{O}$ at nCDN i
- $\mathcal{R}_i = \bigcup_{j \in \mathcal{N}(i)} A_j \rightarrow$ content available from connected nCDNs
- $\beta_i^o(A_{-i}) \triangleq \min_{j \in \mathcal{N}(i)} \{\beta_i^j | o \in A_j\} \rightarrow$ lowest latency to retrieve item o

Average latency experienced by customers of operator i :

$$C_i(A_i, A_{-i}) = \sum_{A_i} w_i^o \alpha_i + \sum_{\mathcal{R}_i} w_i^o \beta_i^o(A_{-i})$$

Interconnected nCDNs - The Cost Model

- K_i storage capacity of nCDN i
- $w_i^o \in \mathbb{R}_+$ demand for item $o \in \mathcal{O}$ at nCDN i
- $\mathcal{R}_i = \bigcup_{j \in \mathcal{N}(i)} A_j \rightarrow$ content available from connected nCDNs
- $\beta_i^o(A_{-i}) \triangleq \min_{j \in \mathcal{N}(i)} \{\beta_i^j | o \in A_j\} \rightarrow$ lowest latency to retrieve item o

Average latency experienced by customers of operator i :

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

Objective

Distributed algorithm - desiderata

1. nCDN i exchange information only with connected $\mathcal{N}(i)$

Objective

Distributed algorithm - desiderata

1. nCDN i exchange information only with connected $\mathcal{N}(i)$
2. Reveals little information about w_i^o

Objective

Distributed algorithm - desiderata

1. nCDN i exchange information only with connected $\mathcal{N}(i)$
2. Reveals little information about w_i^o
3. Leads to content allocation $\bar{A} = (A_i)_{i \in N}$ that is *individually rational*

Objective

Distributed algorithm - desiderata

1. nCDN i exchange information only with connected $\mathcal{N}(i)$
2. Reveals little information about w_i^o
3. Leads to content allocation $\bar{A} = (A_i)_{i \in N}$ that is *individually rational*

Individual rationality

- $r_i(A)$ cost saving ratio for nCDN i in allocation A

$$r_i(A) = \frac{\text{cost saving coop.}}{\text{cost saving } \textbf{no} \text{ coop.}}$$

Objective

Distributed algorithm - desiderata

1. nCDN i exchange information only with connected $\mathcal{N}(i)$
2. Reveals little information about w_i^o
3. Leads to content allocation $\bar{A} = (A_i)_{i \in N}$ that is *individually rational*

Individual rationality

- $r_i(A)$ cost saving ratio for nCDN i in allocation A

$$r_i(A) = \frac{\text{cost saving coop.}}{\text{cost saving } \textbf{no} \text{ coop.}}$$

$r_i(A) > 1$

nCDN i benefits from cooperation

Objective

Distributed algorithm - desiderata

1. nCDN i exchange information only with connected $\mathcal{N}(i)$
2. Reveals little information about w_i^o
3. Leads to content allocation $\bar{A} = (A_i)_{i \in N}$ that is *individually rational*

Individual rationality

- $r_i(A)$ cost saving ratio for nCDN i in allocation A

$$r_i(A) = \frac{\text{cost saving coop.}}{\text{cost saving } \textbf{no} \text{ coop.}}$$

$r_i(A) > 1$
nCDN i benefits from cooperation

$r_i(A) \geq 1$
A_i **individually rational**

Self-enforcing Content Allocations

- Every nCDN i allocates A_i to minimize $C_i(A_i, A_{-i})$ given A_{-i}
- Compatible with operators' selfish interests

Self-enforcing Content Allocations

- Every nCDN i allocates A_i to minimize $C_i(A_i, A_{-i})$ given A_{-i}
- Compatible with operators' selfish interests
- Enforced without bilateral payments $\rightarrow$ self-enforcing allocation

$$\boxed{(A_i^*)_{i\in N} \quad \text{s.t.} \quad A_i^* \in \arg\min_{A_i} C_i(A_i, A_{-i}^*).}$$

$$\Downarrow$$

Nash Equilibrium of strategic game $\Gamma =< N, (\mathcal{A}_i)_{i\in N}, (C_i)_{i\in N} >$

Self-enforcing Content Allocations

- Every nCDN i allocates A_i to minimize $C_i(A_i, A_{-i})$ given A_{-i}
- Compatible with operators' selfish interests
- Enforced without bilateral payments $\rightarrow$ self-enforcing allocation

$$\boxed{(A_i^*)_{i \in N} \quad \text{s.t.} \quad A_i^* \in \arg\min_{A_i} C_i(A_i, A_{-i}^*).}$$

$$\Downarrow$$

Nash Equilibrium of strategic game $\Gamma =< N, (\mathcal{A}_i)_{i \in N}, (C_i)_{i \in N} >$

Distributed *Local-Greedy* Algorithm

- For each nCDN i
- Compute cost saving for each $o \in \mathcal{O}$ given A_{-i}

$$CS_i^o(1, A_{-i}) = \begin{cases} w_i^o \left[\gamma_i - \alpha_i\right] & \text{if } o \notin \mathcal{R}_i \\ w_i^o \left[\beta_i^o(A_{-i}) - \alpha_i\right] & \text{if } o \in \mathcal{R}_i \end{cases}$$

- Store K_i items with highest cost saving

Failure of Distributed *Local-Greedy*

Open questions

- $\exists$ self-enforcing content allocation $A^* = (A_i^*)_{i \in N}$
- Convergence of Distributed *Local-Greedy* to A^*

Failure of Distributed *Local-Greedy*

Open questions

- $\exists$ self-enforcing content allocation $A^* = (A_i^*)_{i \in N}$
- Convergence of Distributed *Local-Greedy* to A^*

Example

- nCDNs $N = \{1, \ldots, 5\}$
- Content items $\mathcal{O} = \{a, b, c, d\}$
- $\exists$ α_i, β_i^j, γ_i and w_i^o

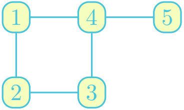

Failure of Distributed *Local-Greedy*

Open questions

- $\exists$ self-enforcing content allocation $A^* = (A_i^*)_{i \in N}$
- Convergence of Distributed *Local-Greedy* to A^*

Example

- nCDNs $N = \{1, \ldots, 5\}$
- Content items $\mathcal{O} = \{a, b, c, d\}$
- $\exists$ α_i, β_i^j, γ_i and w_i^o

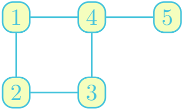

$$(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}) \underset{4}{\rightarrow} (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{b}) \underset{3}{\rightarrow} (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{b}) \underset{2}{\rightarrow} (\boldsymbol{a}, \boldsymbol{c}, \boldsymbol{d}, \boldsymbol{b}) \underset{1}{\rightarrow} (b, c, d, b)$$

$$\underset{4}{\rightarrow} (b, c, d, d) \underset{3}{\rightarrow} (b, c, c, d) \underset{2}{\rightarrow} (b, b, c, d) \underset{1}{\rightarrow} (a, b, c, d)$$

Failure of Distributed *Local-Greedy*

Open questions

- $\exists$ self-enforcing content allocation $A^* = (A_i^*)_{i \in N}$
- Convergence of Distributed *Local-Greedy* to A^*

Example

- nCDNs $N = \{1, \ldots, 5\}$
- Content items $\mathcal{O} = \{a, b, c, d\}$
- $\exists$ α_i, β_i^j, γ_i and w_i^o

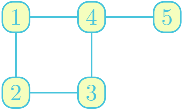

$$(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}) \underset{4}{\rightarrow} (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{b}) \underset{3}{\rightarrow} (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{b}) \underset{2}{\rightarrow} (\boldsymbol{a}, \boldsymbol{c}, \boldsymbol{d}, \boldsymbol{b}) \underset{1}{\rightarrow} (b, c, d, b)$$

$$\underset{4}{\rightarrow} (b, c, d, d) \underset{3}{\rightarrow} (b, c, c, d) \underset{2}{\rightarrow} (b, b, c, d) \underset{1}{\rightarrow} (a, b, c, d)$$

Theorem: Self-enforcing content allocations might not exist

Failure of Distributed *Local-Greedy*

Open questions

- $\exists$ self-enforcing content allocation $A^* = (A_i^*)_{i \in N}$
- Convergence of Distributed *Local-Greedy* to A^*

Example

- nCDNs $N = \{1, \dots, 5\}$
- Content items $\mathcal{O} = \{a, b, c, d\}$
- $\exists\ \alpha_i,\ \beta_i^j,\ \gamma_i$ and w_i^o

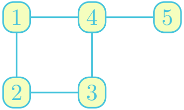

$$(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}) \underset{4}{\rightarrow} (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{b}) \underset{3}{\rightarrow} (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{b}) \underset{2}{\rightarrow} (\boldsymbol{a}, \boldsymbol{c}, \boldsymbol{d}, \boldsymbol{b}) \underset{1}{\rightarrow} (b, c, d, b)$$

$$\underset{4}{\rightarrow} (b, c, d, d) \underset{3}{\rightarrow} (b, c, c, d) \underset{2}{\rightarrow} (b, b, c, d) \underset{1}{\rightarrow} (a, b, c, d)$$

Theorem: Self-enforcing content allocations might not exist

Corollary: $\alpha_i = \alpha, \gamma_i = \gamma$, and $\beta_i^j = \beta_j^i$ do not help!

Bilateral Compensation-based Algorithms

- Introduce periodic **bilateral payments** among nCDNs
- **Open question**: $\exists$ stable content allocation A

Bilateral Compensation-based Algorithms

- Introduce periodic **bilateral payments** among nCDNs
- **Open question**: $\exists$ stable content allocation A

Bilateral compensation-based algorithms in a nutshell

- At every time step t
  1. nCDNs in set N_t allowed to propose allocation update
  2. Connected nCDNs $j \in \mathcal{N}(N_t)$ can offer payments to dissuade N_t

Bilateral Compensation-based Algorithms

- Introduce periodic **bilateral payments** among nCDNs
- **Open question**: $\exists$ stable content allocation A

Bilateral compensation-based algorithms in a nutshell

- At every time step t
  1. nCDNs in set N_t allowed to propose allocation update
  2. Connected nCDNs $j \in \mathcal{N}(N_t)$ can offer payments to dissuade N_t

Synchronization schemes

Asynchronous	**Plesiochronous**	**Synchronous**
$\lvert N_t \rvert = 1$	$N_t \subset N$	$N_t = N$

Aggregate-value Compensation Algorithm (AC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$

Aggregate-value Compensation Algorithm (AC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$

Aggregate-value Compensation Algorithm (AC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. All nCDNs $j \in \mathcal{N}(i)$ s.t. $\Delta C_j > 0$ offer a compensation $p_j^i = \Delta C_j$

Aggregate-value Compensation Algorithm (AC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. All nCDNs $j \in \mathcal{N}(i)$ s.t. $\Delta C_j > 0$ offer a compensation $p_j^i = \Delta C_j$
  4. nCDN i updates its content allocation to $A_i(t)$ **only if**

$$\sum_{j \in \mathcal{N}(i)} p_j^i < -\Delta C_i(t)$$

Aggregate-value Compensation Algorithm (AC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. All nCDNs $j \in \mathcal{N}(i)$ s.t. $\Delta C_j > 0$ offer a compensation $p_j^i = \Delta C_j$
  4. nCDN i updates its content allocation to $A_i(t)$ **only if**

$$\sum_{j \in \mathcal{N}(i)} p_j^i < -\Delta C_i(t)$$

Asynchronous operation, $|N_t| = 1$

Theorem: AC algorithm terminates in a finite number of steps (1-AC)

Aggregate-value Compensation Algorithm (AC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. All nCDNs $j \in \mathcal{N}(i)$ s.t. $\Delta C_j > 0$ offer a compensation $p_j^i = \Delta C_j$
  4. nCDN i updates its content allocation to $A_i(t)$ **only if**

$$\sum_{j \in \mathcal{N}(i)} p_j^i < -\Delta C_i(t)$$

Asynchronous operation, $|N_t| = 1$

Theorem: AC algorithm terminates in a finite number of steps (1-AC)

Plesiochronous operation, $N_t \subset N$

Theorem: If N_t is 2-independent set of $\mathcal{G} \Rightarrow$ AC algorithm terminates in a finite number of steps ($\mathcal{I}^2$-AC)

- Distance between any two vertexes of N_t is at least 3
  - Few nodes update at each time step
  - Needs coordination scheme over 2-hops neighbours

Object-value Compensation Algorithm (OC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$

Object-value Compensation Algorithm (OC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. $j \in \mathcal{N}(N_t)$ offers compensation $p^k_{j,o} = \Delta C^o_j$ **for each individual item** o to nCDN $k \in N_t$ s.t

Object-value Compensation Algorithm (OC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. $j \in \mathcal{N}(N_t)$ offers compensation $p^k_{j,o} = \Delta C^o_j$ **for each individual item** o to nCDN $k \in N_t$ s.t
     - nCDN k evicting item o **and**
     - lowest latency, i.e. $k = \min_{i \in \mathcal{N}(j)} \{\beta^i_j | o \in A_i\}$

Object-value Compensation Algorithm (OC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. $j \in \mathcal{N}(N_t)$ offers compensation $p^k_{j,o} = \Delta C^o_j$ **for each individual item** o to nCDN $k \in N_t$ s.t
     - nCDN k evicting item o **and**
     - lowest latency, i.e. $k = \min_{i \in \mathcal{N}(j)} \{\beta^i_j | o \in A_i\}$
  4. nCDN i updates its content allocation to $A_i(t)$ **only if**

$$\sum_{j \in \mathcal{N}(i)} \sum_{o \in \mathcal{O}} p^i_{j,o} < -\Delta C_i(t)$$

Object-value Compensation Algorithm (OC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. $j \in \mathcal{N}(N_t)$ offers compensation $p^k_{j,o} = \Delta C^o_j$ **for each individual item** o to nCDN $k \in N_t$ s.t
     - nCDN k evicting item o **and**
     - lowest latency, i.e. $k = \min_{i \in \mathcal{N}(j)} \{\beta^i_j | o \in A_i\}$
  4. nCDN i updates its content allocation to $A_i(t)$ **only if**

$$\sum_{j \in \mathcal{N}(i)} \sum_{o \in \mathcal{O}} p^i_{j,o} < -\Delta C_i(t)$$

Plesiochronous operation, $N_t \subset N$

Theorem: If N_t is 1-independent set of $\mathcal{G} \Rightarrow$ OC algorithm terminates in a finite number of steps ($\mathcal{I}^1$-OC)

Object-value Compensation Algorithm (OC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. $j \in \mathcal{N}(N_t)$ offers compensation $p^k_{j,o} = \Delta C^o_j$ **for each individual item** o to nCDN $k \in N_t$ s.t
     - nCDN k evicting item o **and**
     - lowest latency, i.e. $k = \min_{i \in \mathcal{N}(j)}\{\beta^i_j | o \in A_i\}$
  4. nCDN i updates its content allocation to $A_i(t)$ **only if**

$$\sum_{j \in \mathcal{N}(i)} \sum_{o \in \mathcal{O}} p^i_{j,o} < -\Delta C_i(t)$$

Plesiochronous operation, $N_t \subset N$

Theorem: If N_t is 1-independent set of $\mathcal{G} \Rightarrow$ OC algorithm terminates in a finite number of steps ($\mathcal{I}^1$-OC)

- $\uparrow$ On average $|\mathcal{I}^1| >> |\mathcal{I}^2|$
- $\downarrow$ $\mathcal{I}^1$-OC reveals more information about w^o_i

Object-value Compensation Algorithm (OC)

- At every time step t
  1. $\forall$ nCDN $i \in N_t$ computes $A_i(t)$ that decreases its cost, i.e. $\Delta C_i < 0$
  2. nCDN i communicates $A_i(t)$ to all neighboring nCDNs $j \in \mathcal{N}(i)$
  3. $j \in \mathcal{N}(N_t)$ offers compensation $p^k_{j,o} = \Delta C^o_j$ **for each individual item** o to nCDN $k \in N_t$ s.t
     - nCDN k evicting item o **and**
     - lowest latency, i.e. $k = \min_{i \in \mathcal{N}(j)} \{\beta^i_j | o \in A_i\}$
  4. nCDN i updates its content allocation to $A_i(t)$ **only if**

$$\sum_{j \in \mathcal{N}(i)} \sum_{o \in \mathcal{O}} p^i_{j,o} < -\Delta C_i(t)$$

Plesiochronous operation, $N_t \subset N$

Theorem: If N_t is 1-independent set of $\mathcal{G} \Rightarrow$ OC algorithm terminates in a finite number of steps ($\mathcal{I}^1$-OC)

- $\uparrow$ On average $|\mathcal{I}^1| >> |\mathcal{I}^2|$
- $\downarrow$ $\mathcal{I}^1$-OC reveals more information about w^o_i
- Neither AC nor OC ensure individual rationality!

Achieving Individual Rationality

Opt Out scheme

- Set $N_{\text{COOP}} \leftarrow N$, until individual rationality achieved **do**
  1. Run algorithm AC or OC $\leftarrow$ termination in $\boldsymbol{A}$
  2. Exlcude nCDNs k s.t. $r_k(\boldsymbol{A}) < 1$,
     - $N_{\text{COOP}} \leftarrow \{i | r_i(\boldsymbol{A}) \geq 1\}$
     - Excluded nCDNs k do not cooperate, i.e. A_k^I

Achieving Individual Rationality

Opt Out scheme

- Set $N_{\text{COOP}} \leftarrow N$, until individual rationality achieved **do**
  1. Run algorithm AC or OC $\leftarrow$ termination in $\boldsymbol{A}$
  2. Exlcude nCDNs k s.t. $r_k(\boldsymbol{A}) < 1$,
     - $N_{\text{COOP}} \leftarrow \{i | r_i(\boldsymbol{A}) \geq 1\}$
     - Excluded nCDNs k do not cooperate, i.e. A_k^I

Numerical Results - Impact of graph topology

- CAIDA dataset - AS-level topology
- European ASes with $> 2^{16}$ alloc. IPs
- **CAIDA**: $|N| = 638$, avg. node degree 10.8
- **CAIDA-ER**: Erdős-Rényi
- **CAIDA-BA**: Barabási-Albert

Achieving Individual Rationality

OPT OUT scheme

- Set $N_{\text{COOP}} \leftarrow N$, until individual rationality achieved **do**
  1. Run algorithm AC or OC $\leftarrow$ termination in $\boldsymbol{A}$
  2. Exlcude nCDNs k s.t. $r_k(\boldsymbol{A}) < 1$,
     - $N_{\text{COOP}} \leftarrow \{i | r_i(\boldsymbol{A}) \geq 1\}$
     - Excluded nCDNs k do not cooperate, i.e. A_k^I

NUMERICAL RESULTS - IMPACT OF GRAPH TOPOLOGY

- CAIDA dataset - AS-level topology
- European ASes with $> 2^{16}$ alloc. IPs
- **CAIDA**: $|N| = 638$, avg. node degree 10.8
- **CAIDA-ER**: Erdős-Rényi
- **CAIDA-BA**: Barabási-Albert

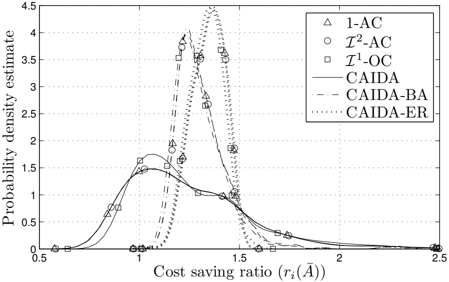

Achieving Individual Rationality

Opt Out scheme

- Set $N_{\text{COOP}} \leftarrow N$, until individual rationality achieved **do**
  1. Run algorithm AC or OC ← termination in $\boldsymbol{A}$
  2. Exlcude nCDNs k s.t. $r_k(\boldsymbol{A}) < 1$,
     - $N_{\text{COOP}} \leftarrow \{i | r_i(\boldsymbol{A}) \geq 1\}$
     - Excluded nCDNs k do not cooperate, i.e. A_k^I

Numerical Results - Impact of graph topology

- CAIDA dataset - AS-level topology
- European ASes with $> 2^{16}$ alloc. IPs
- **CAIDA**: $|N| = 638$, avg. node degree 10.8
- **CAIDA-ER**: Erdős-Rényi
- **CAIDA-BA**: Barabási-Albert

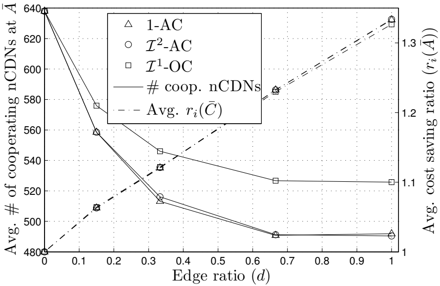

Numerical Results - Convergence Rate

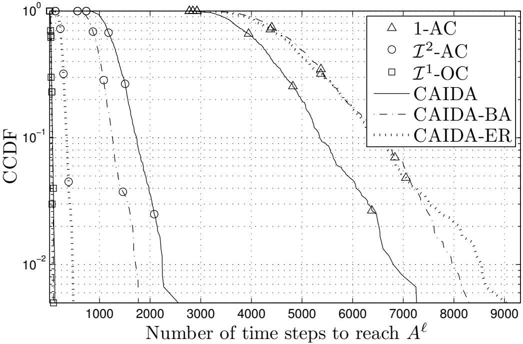

Numerical Results - Convergence Rate

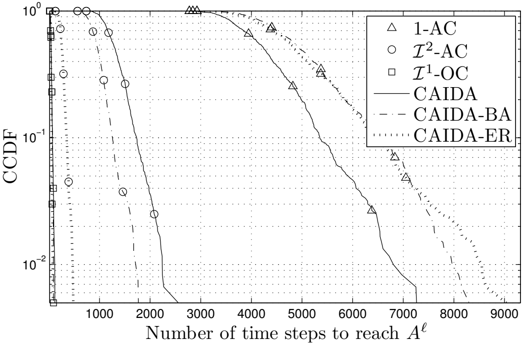

Graph	$\mathcal{I}^1$ sets avg. size	$\mathcal{I}^2$ sets avg. size
CAIDA	39.8	2.9
CAIDA-BA	63.8	4.9
CAIDA-ER	79.8	17.8

Numerical Results - Scaling for Storage Capacity

- Increasing storage capacity K_i at every nCDN $i \in N$

Numerical Results - Scaling for Storage Capacity

- Increasing storage capacity K_i at every nCDN $i \in N$

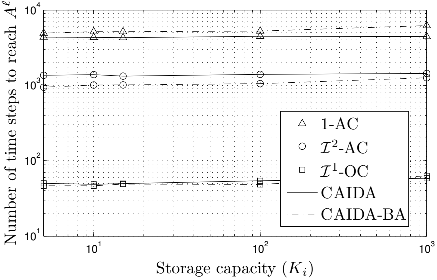

Numerical Results - Scaling for Storage Capacity

- Increasing storage capacity K_i at every nCDN $i \in N$

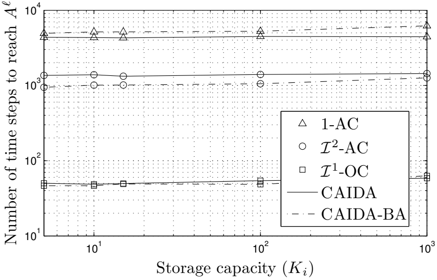

- Convergence rate insensitive to K_i

Numerical Results - Scaling for Storage Capacity

- Increasing storage capacity K_i at every nCDN $i \in N$

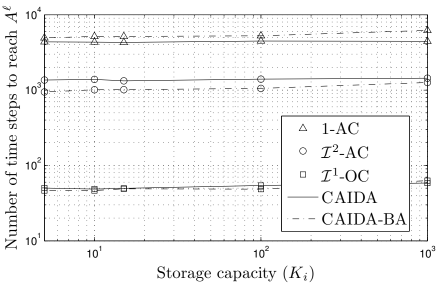

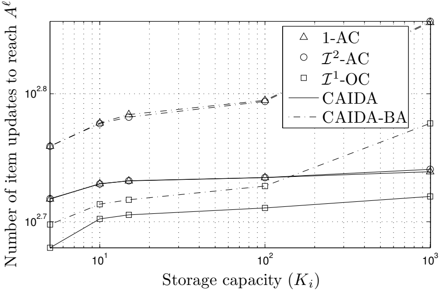

- Convergence rate insensitive to K_i

Numerical Results - Scaling for Storage Capacity

- Increasing storage capacity K_i at every nCDN $i \in N$

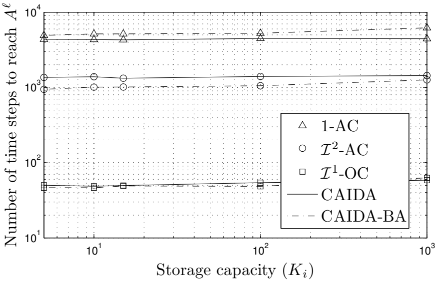

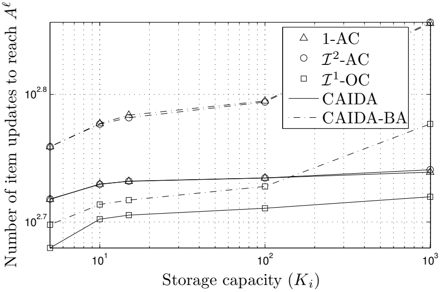

- Convergence rate insensitive to K_i

- $*$-AC: Same # of item updates
- $\mathcal{I}^1$-OC more efficient
- $\uparrow$ # item updates $\neq \uparrow$ conv. rate

Conclusion

Interconnected Operator Managed CDNs

- Retrieval of content from connected nCDNs
- Individually rational content allocation

Conclusion

Interconnected Operator Managed CDNs

- Retrieval of content from connected nCDNs
- Individually rational content allocation

Self-enforcing Content Allocations

- Distributed greedy cost minimization not suitable

Conclusion

Interconnected Operator Managed CDNs

- Retrieval of content from connected nCDNs
- Individually rational content allocation

Self-enforcing Content Allocations

- Distributed greedy cost minimization not suitable

Stable Allocations

- Bilateral compensation-based algorithms reach stable allocations
- Faster convergence at the cost of revealing information

Conclusion

Interconnected Operator Managed CDNs

- Retrieval of content from connected nCDNs
- Individually rational content allocation

Self-enforcing Content Allocations

- Distributed greedy cost minimization not suitable

Stable Allocations

- Bilateral compensation-based algorithms reach stable allocations
- Faster convergence at the cost of revealing information

Individual Rationality

- Individual rationality of stable allocations depends on graph topology
- At least 80% of nCDNs have incentive to cooperate

Distributed Algorithms for Content Allocation in Interconnected Content Distribution Networks

Valentino Pacifici, György Dán

Laboratory for Communication Networks
KTH Royal Institute of Technology
Stockholm, Sweden

Hong Kong, April 30, 2015