Native Content Distribution through Off-Path Content Discovery

A Proposal for a "Downstream FIB"

"Opportunistic Off-Path Content Discovery in Information-Centric Networks"
O. Ascigil, V. Sourlas, I. Psaras, G. Pavlou
IEEE LANMAN 2016
Best Paper Award

Ioannis Psaras
EPSRC Fellow
University College London
`i.psaras@ucl.ac.uk`

"Information Resilience Through User-Assisted Caching in Disruptive Content-Centric Networks"
V. Sourlas, L. Tassiulas, I. Psaras, G. Pavlou
IFIP NETWORKING 2015
Best Paper Award

ICN Promise

Transform the Internet to a
Native Content Distribution Network

1. Name content
2. Route on names – stateful forwarding
3. Enable and exploit in-network caching
4. Find nearest copy of content *in on-path caches*!

Is the goal achieved?

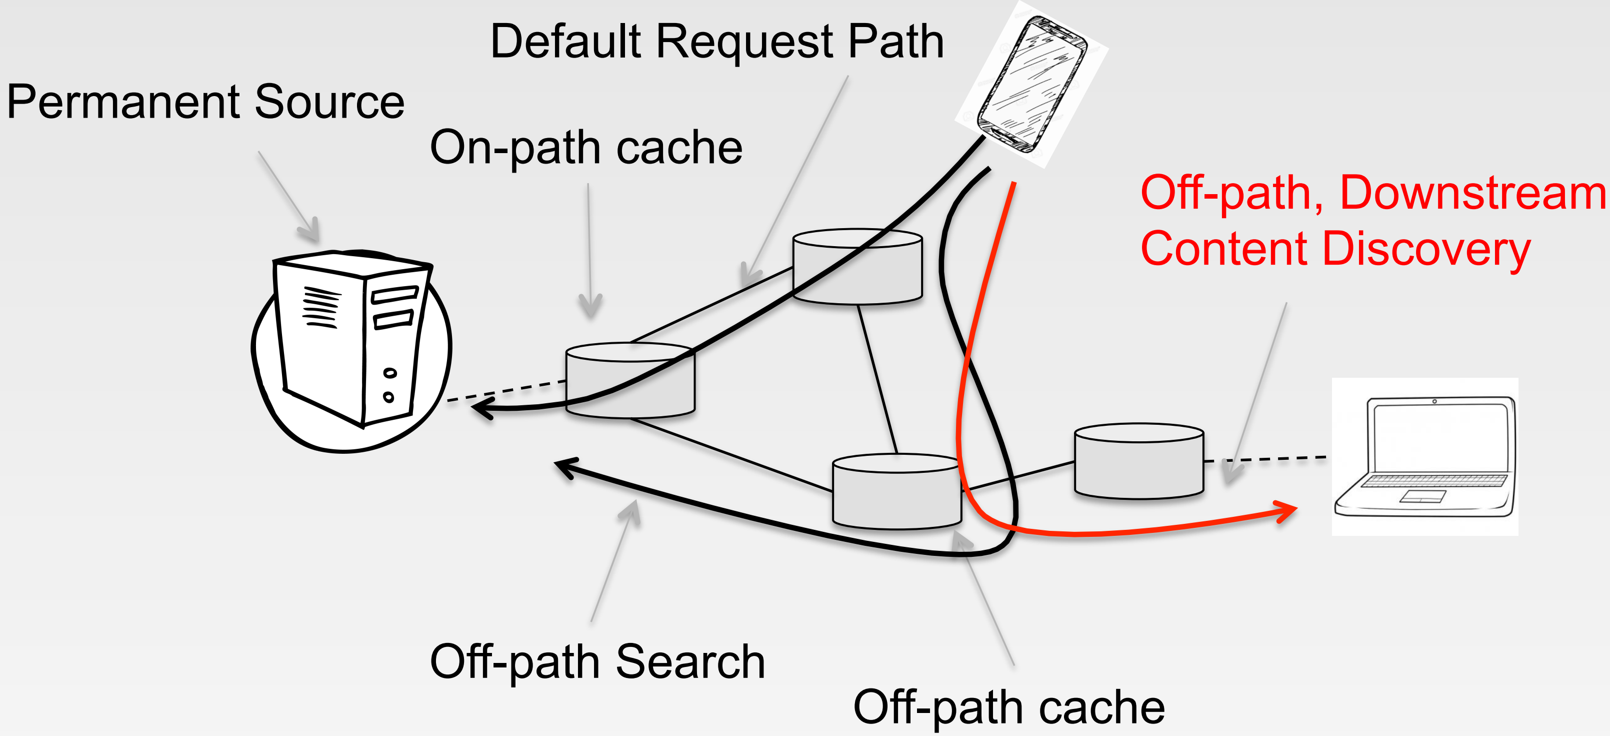

There is always a permanent source node

Requests/Interests always follow breadcrumbs towards the source node – through FIB

Off-path caching mechanisms attempt to find content in the vicinity – significant overhead introduced

There is no mechanism to point to alternative sources, *e.g.*, sources that have recently requested the content

Opportunistic Content Discovery

A Proposal for "Downstream FIB"

Stateful forwarding of data packets: *data packets leave breadcrumbs*

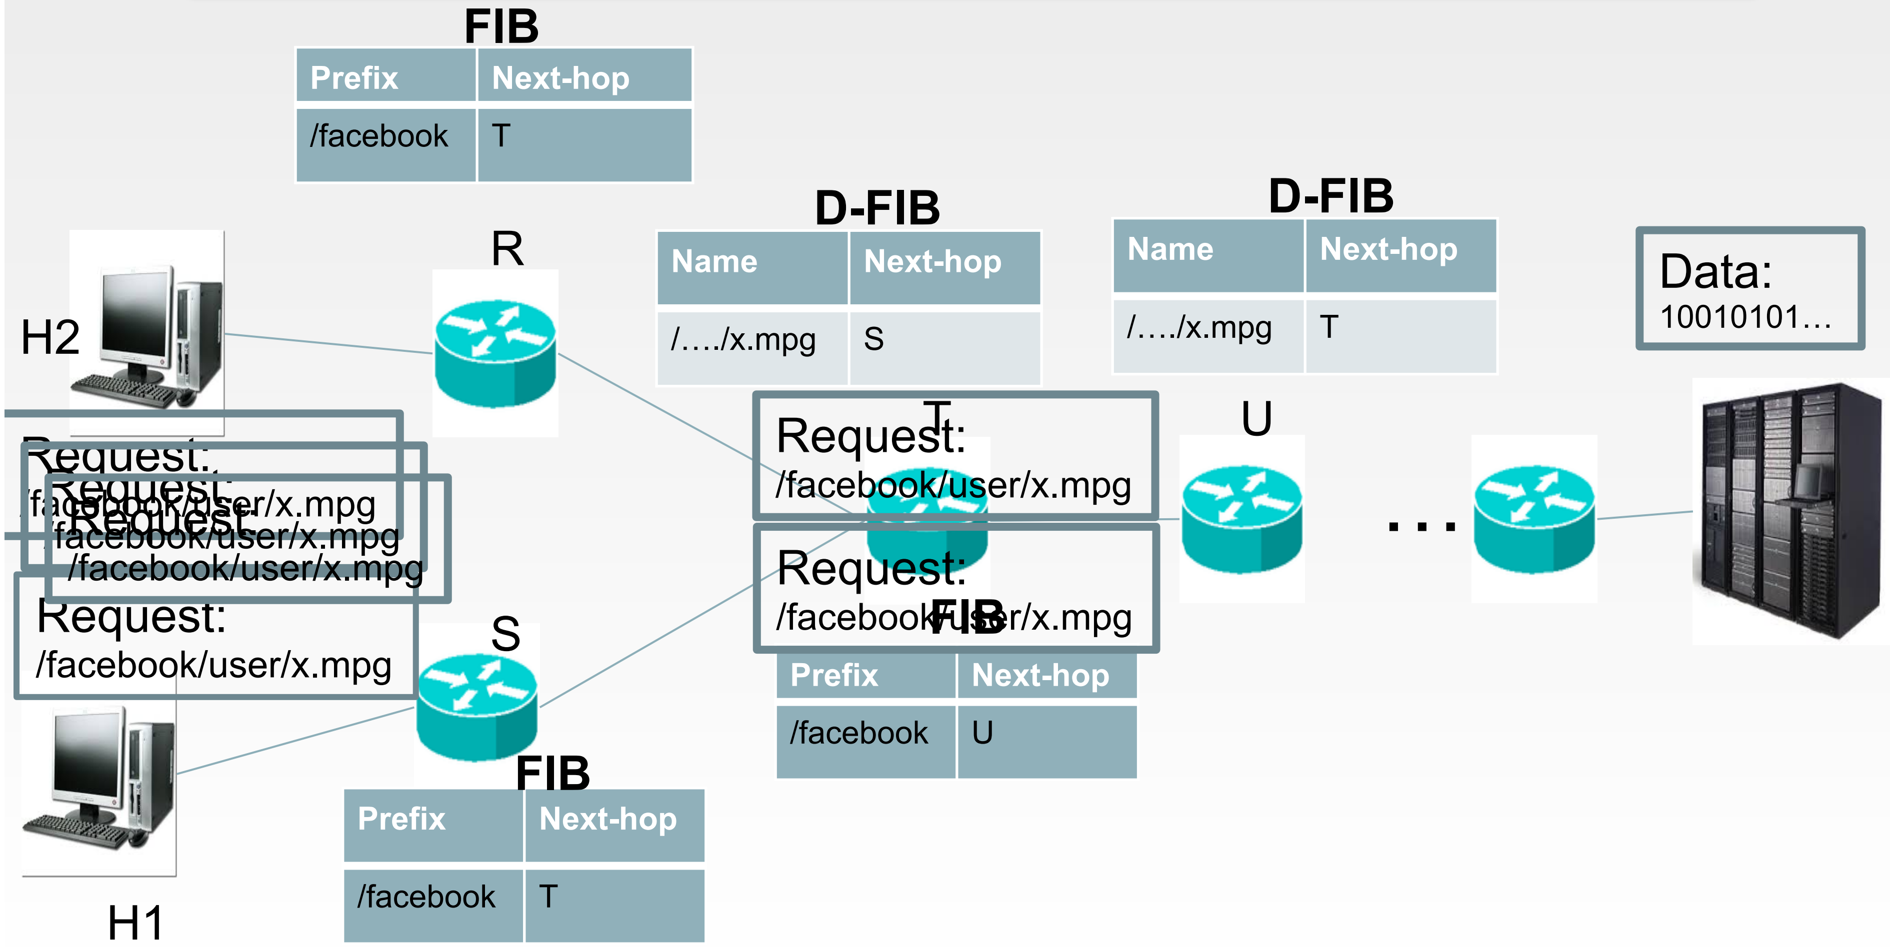

Opportunistic Content Discovery: Downstream FIB Table

- Content Store (CS)
- Pending Interest Table (PIT)
- Forwarding Information Base (FIB)

Same to NDN original model

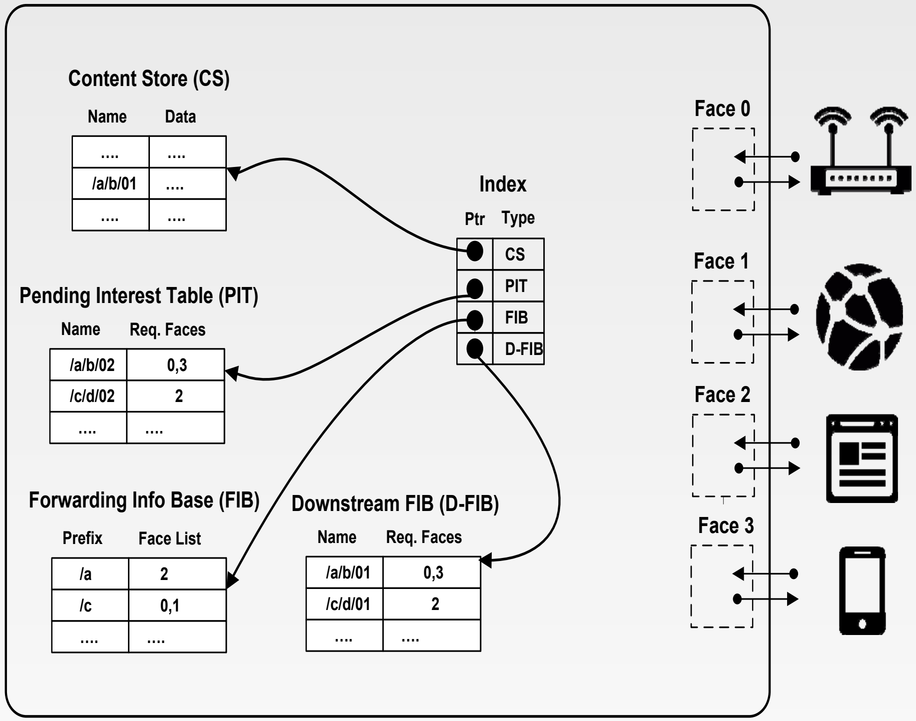

Downstream FIB (D-FIB)

- Keeps track of data packet next hop.
- "*Breadcrumbs*" for user-assisted caching.
- Allows for a list of outgoing faces.
- Similar to Persistent Interests (PI) in C. Tsilopoulos and G. Xylomenos, "Supporting Diverse Traffic Types in ICN" ACM SIGCOMM ICN 2011.

Opportunistic Content Discovery: Routing using D-FIB & FIB

- Goal:
  - Introduce alternative content sources, not towards the original source
  - **limit overhead and reduce the number of requests reaching the content origin**
- Expected Results:
  - Increase Cache Hits (downstream)
  - Reduce delivery latency (number of hops traveled)
- Challenge:
  - How do we manage incoming interests
  - Which path should requests follow:
    - Upstream
    - Downstream
    - Or both..

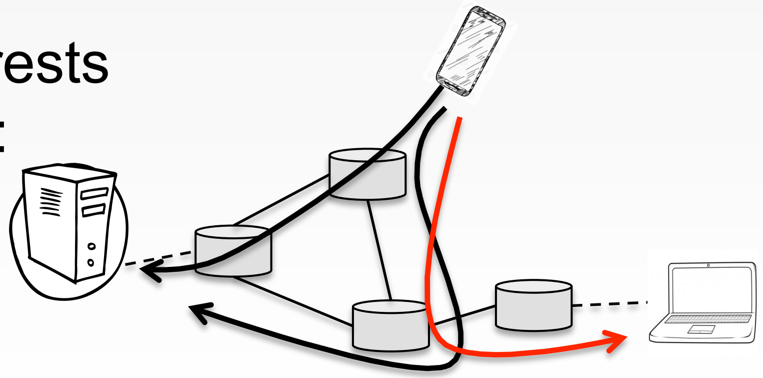

Opportunistic Content Discovery:

Addressing the request management challenge

- Each request is associated with a ***Total Forwarding Counter (TFC) value***
  - spend it on sending a copy of a request *downstream*
  - spend it on following the FIB table towards the content origin (*upstream*)
  - spend it on both (*multicast*)
- TFC is initially set by the access router
- New Forwarding Strategies based on D-FIB
  - Determines how **TFC quota** is spent at each router

Downstream FIB Table

The **Multicast** Case

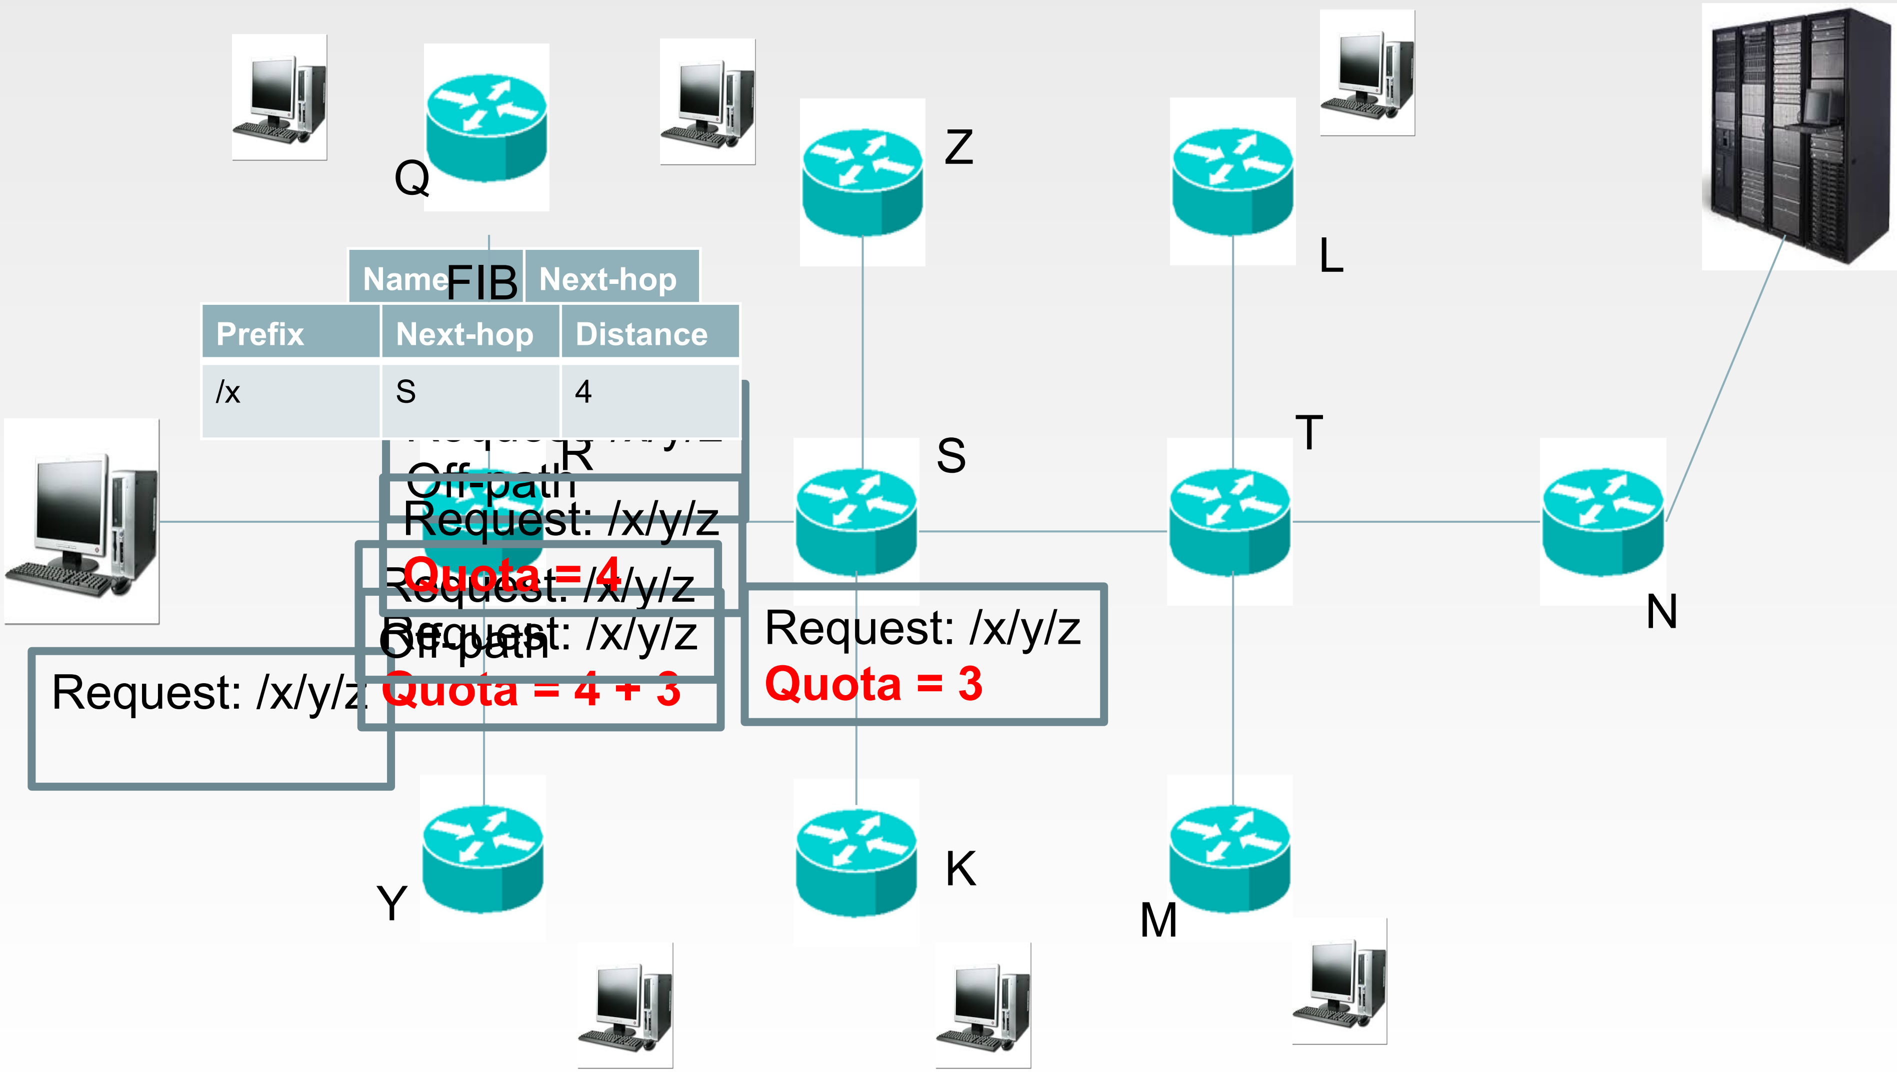

Prefix	Next-hop	Distance
/x	S	4

Opportunistic Content Discovery: Forwarding Strategies

- Check Content Store; if no matching content, then:
- Lookup FIB and D-FIB
  - If D-FIB returns no entries, follow FIB (forward upstream)
  - If D-FIB returns one or more entries, then the **forwarding strategy** decides what action to perform
- Two simple strategies:
  - **ALL strategy**: Send a copy of the request to all the next-hops in the D-FIB entry
    - the cache is **closer (number of hops)** than the content origin
  - **ONE strategy**: Send a copy of the request to only one next-hop in the D-FIB entry
    - **Freshest** entry which is closer than the content origin

Performance Evaluation

Performance Evaluation Setup

- Implemented our approach in **ndnSIM** — an ns-3 based simulator
- **Performance metrics:**
  - **Cache hit ratio:** percentage of the interests that have been satisfied
    - Off-path/on-path
  - **The minimum hop distance:** number of hops traveled by the (first) data arriving at the user from a responding router or the content origin for each successful request
  - **The mean traffic overhead:** the mean number of hops that the initiated Data packets travel in the network
- **Variables**:
  - Cache size at each node
  - D-FIB size w.r.t. content population size
  - Initial Quota

Performance Evaluation Setup

- **Using a RocketFuel topology: AS 4755 VSNL (India)**
  - 191 nodes: 148 edge, 39 gateway, and 4 backbone routers
  - 242 bi-directional links
  - Average distance from edge-routers to producer: 3.5
- **Request rate: 100 requests/sec**
  - Randomly select an **edge** router
- **Content Population: 10,000**
  - One chunk per item
- **One content server**
  - attached to a randomly chosen edge router
  - our results comparing performance of on-path/off-path is best-case scenario
- **Popularity of the items determined by a Zipf law of exponents**
  - Zipf parameter *z*: 0.7
- **Total Forwarding Counter Quota:** Shortest path length + 3
- **Duration:** 1 hour (following an hour of warm-up phase)

Evaluation: Impact of Router's Cache Size

- Impact of D-FIB size w.r.t. content population on the performance

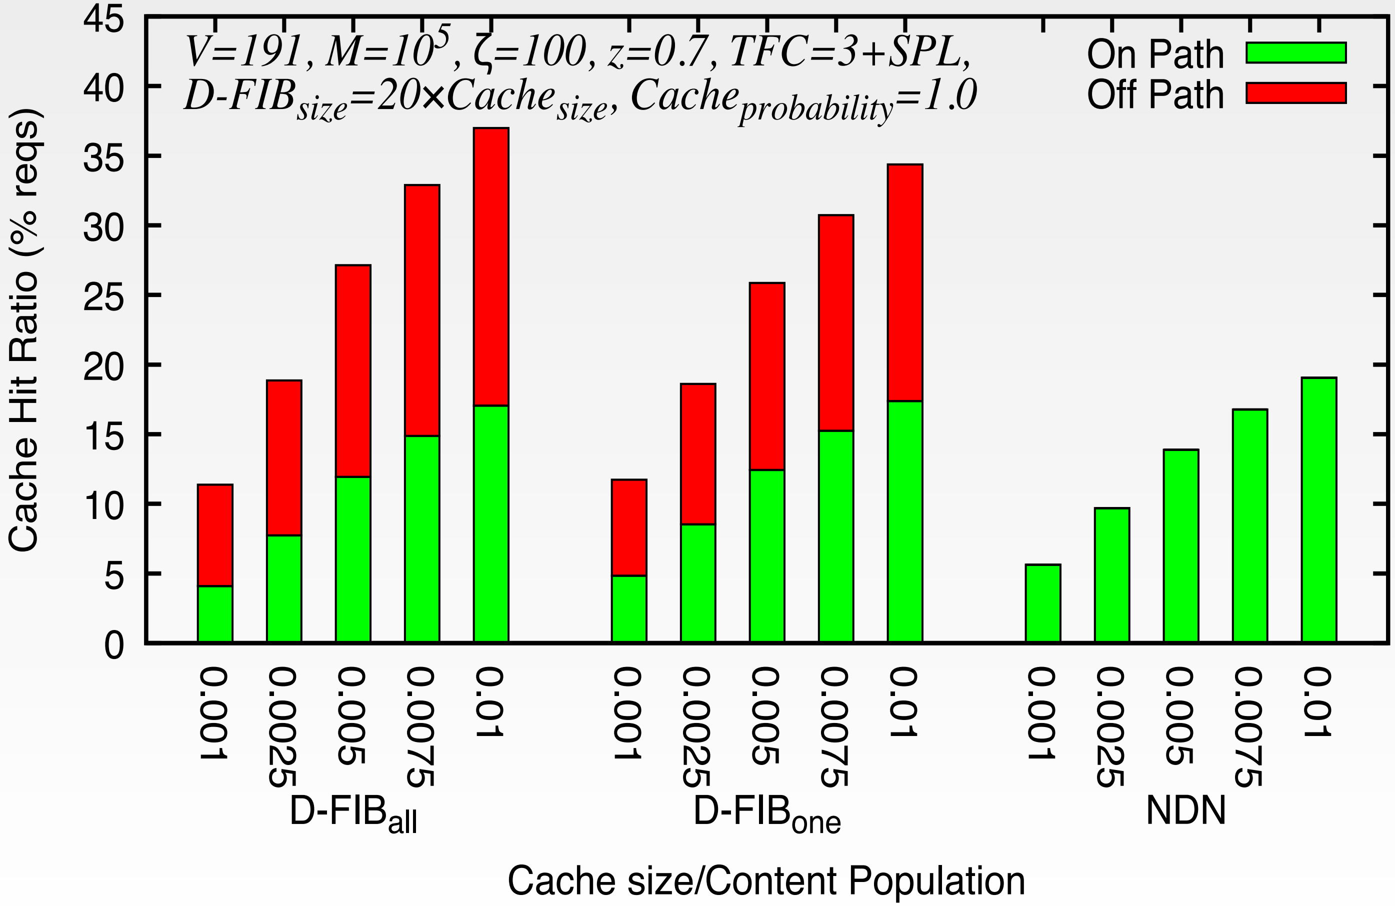

Evaluation: Impact of Router's Cache Size

Average edge-router to source hop-distance: 3.5

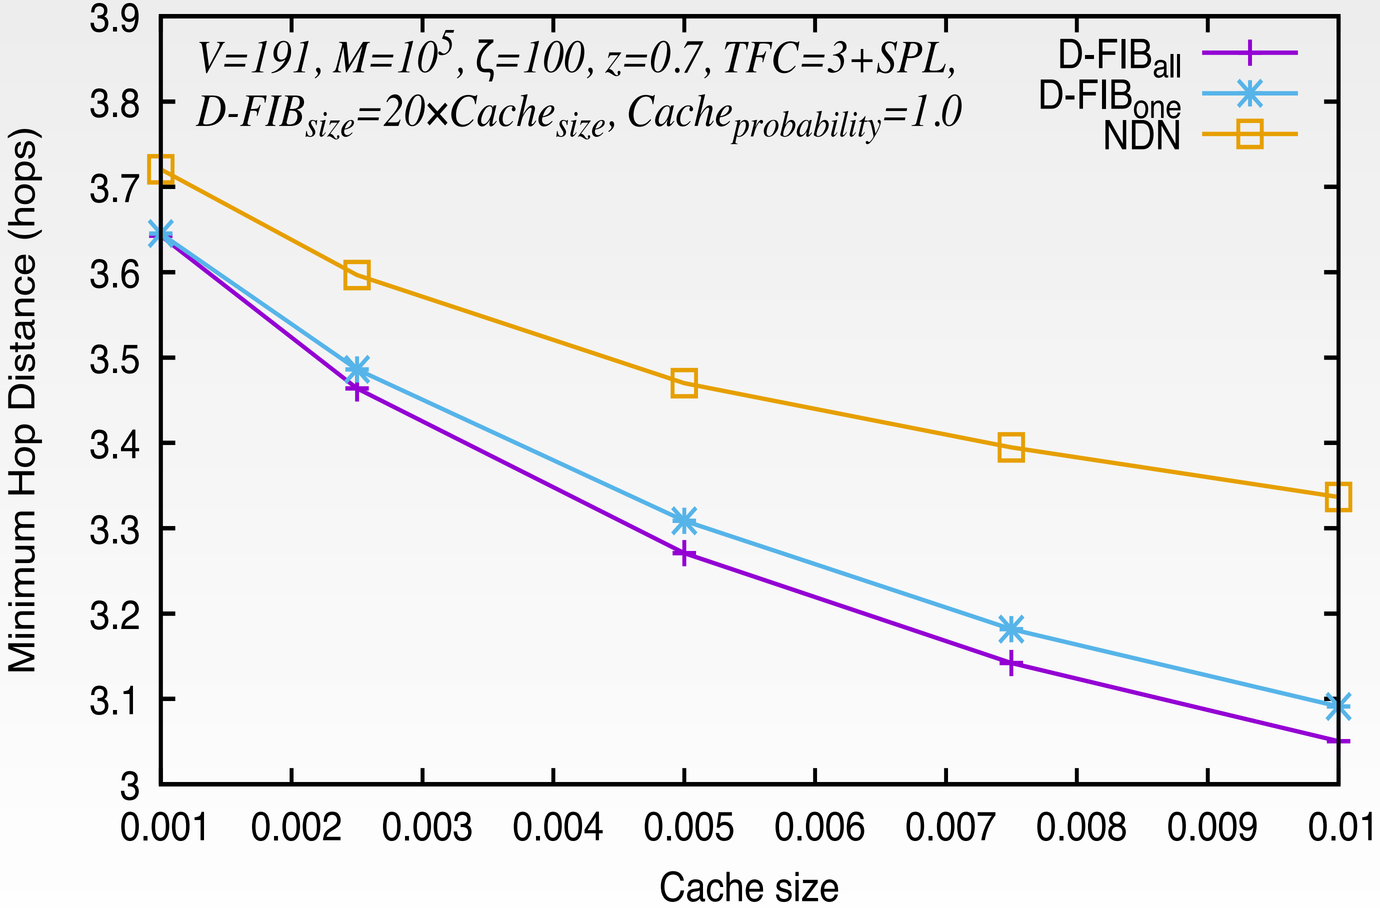

Evaluation: Impact of Router's Cache Size

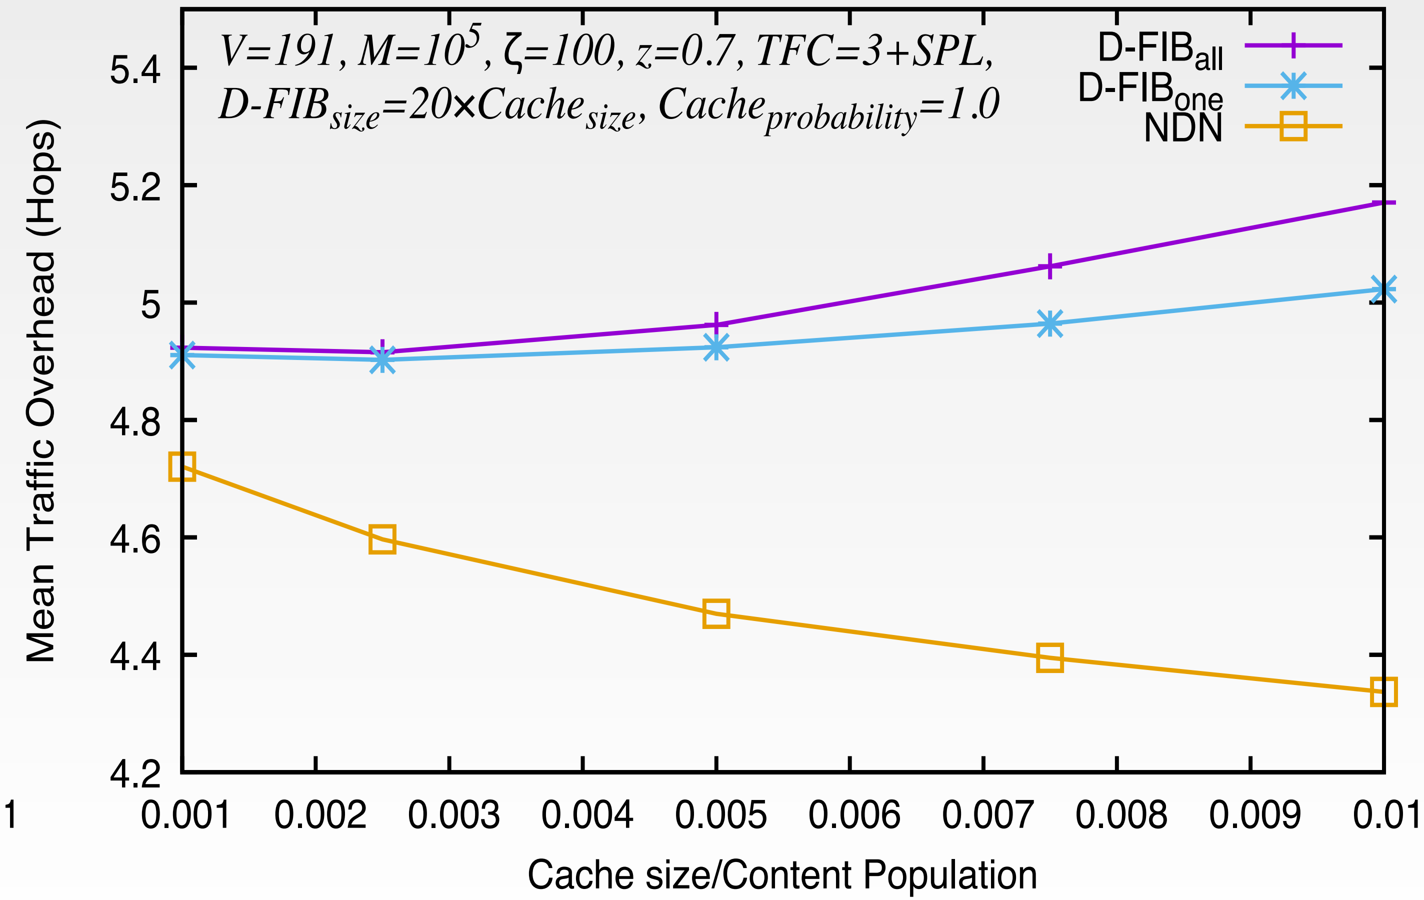

Evaluation: Impact of D-FIB size

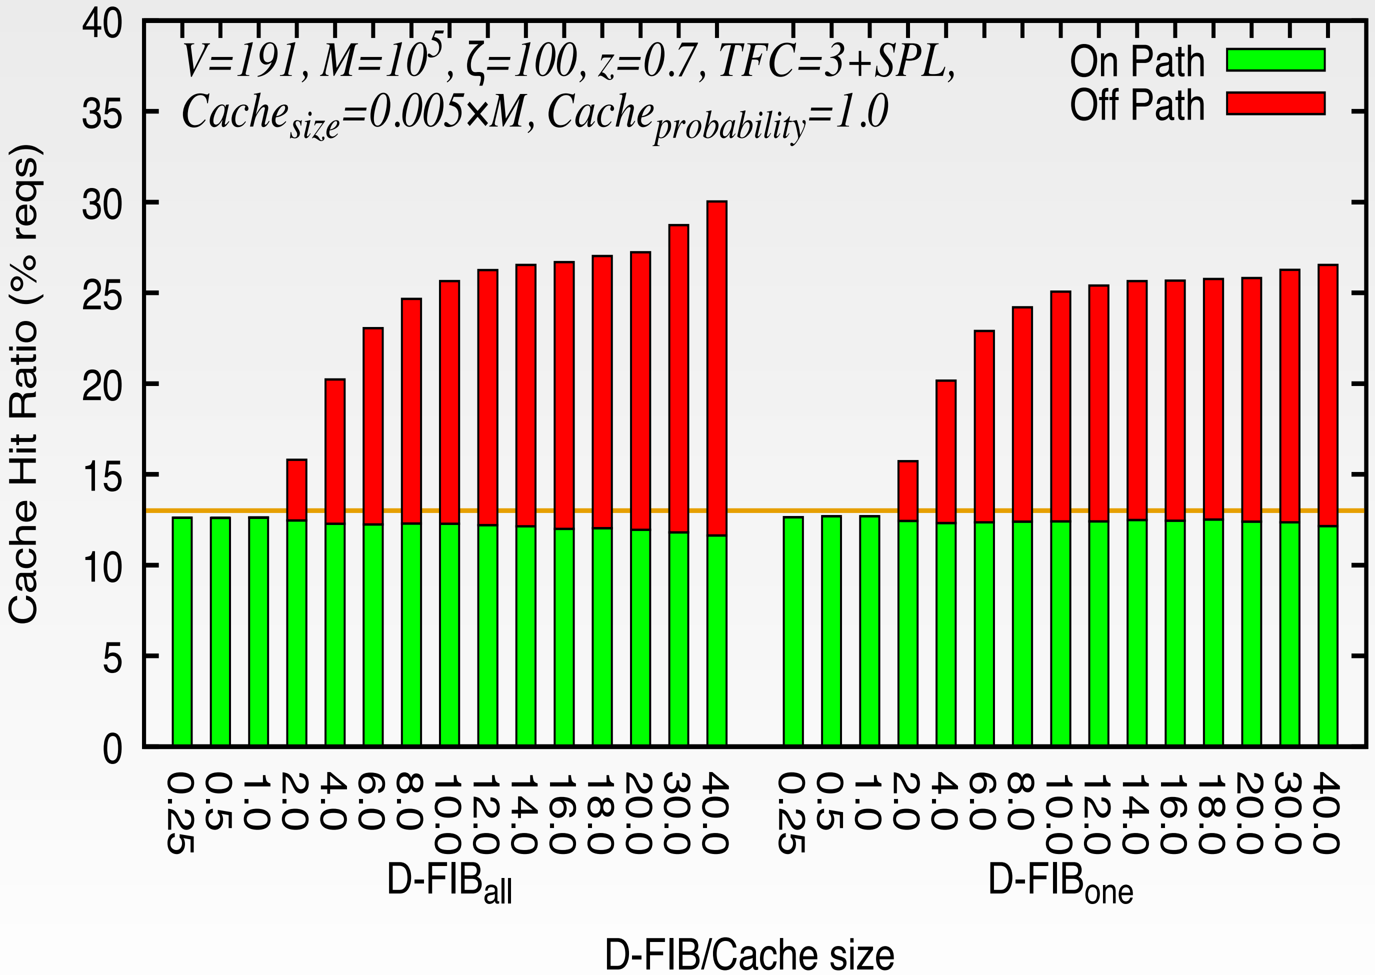

Evaluation: Impact of Initial Quota

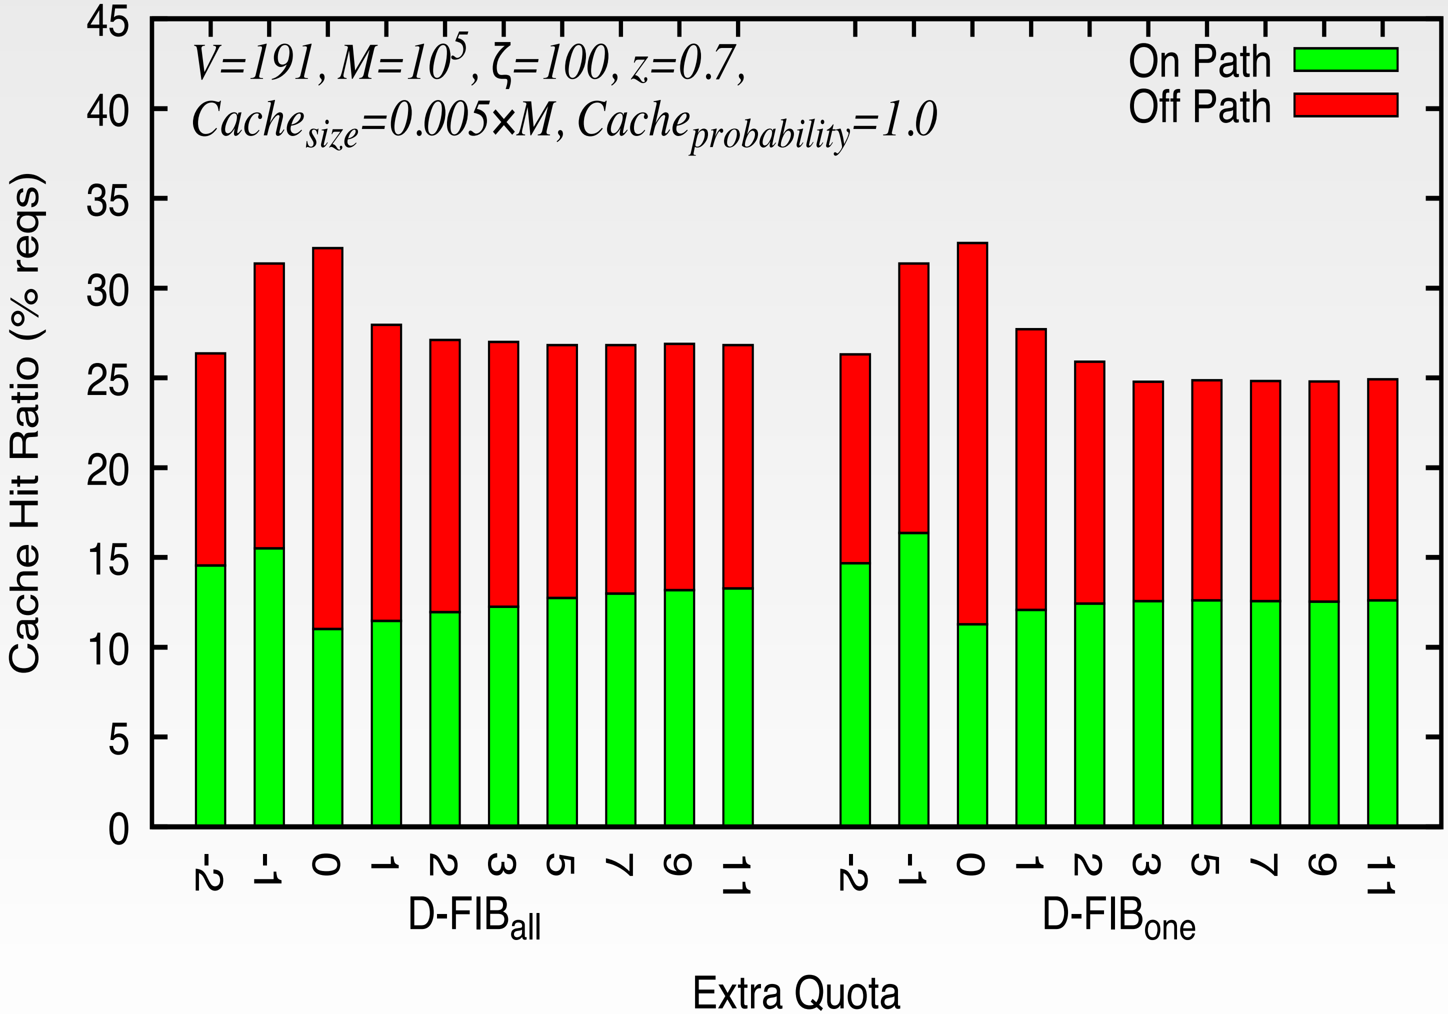

Evaluation: Impact of Initial Quota

What percentage of the first requests manage to fetch the content?

Extra Quota	*Sat.* $Rate_{all}$	*Sat.* $Rate_{one}$
-2	26.3%	26.3%
-1	31.3%	31.3%
0	91.3%	91.3%
1	98.5%	98.5%
2	99.2%	99.7%
3	99.7%	99.9%
...	...	...
11	100%	100%

TABLE I

REQUEST SATISFACTION RATE OF THE FIRST REQUESTS FOR DIFFERENT EXTRA QUOTA VALUES

Opportunistic Off-Path Content Discovery

"Opportunistic Off-Path Content Discovery in Information-Centric Networks"
O. Ascigil, V. Sourlas, I. Psaras, G. Pavlou
IEEE LANMAN 2016
Best Paper Award

Information Resilience through D-FIB in Fragmented Networks

"Information Resilience Through User-Assisted Caching in Disruptive Content-Centric Networks"
V. Sourlas, L. Tassiulas, I. Psaras, G. Pavlou
IFIP NETWORKING 2015
Best Paper Award

Problem Attacked

When the network gets fragmented, and given we have a number of (in-network) caches, for how long can we keep the content “alive” in caches and end-user devices?

- How do we find “alive” content (i.e., content still in caches)?

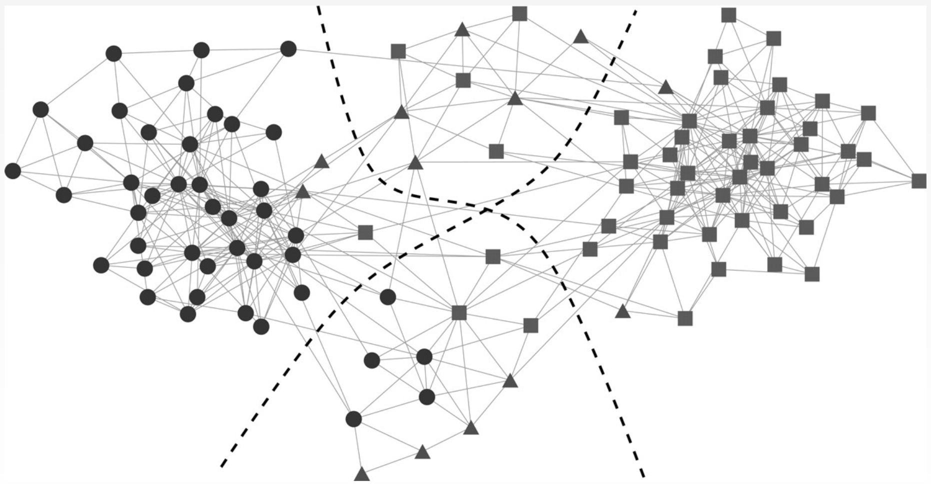

Goals

- Find ways to:
  - Exploit all possible sources to retrieve content when the main path is “down”
  - Exploit in-network caching to prolong information lifetime in case of disasters
  - Natively support P2P-like content distribution at the network layer

Information Resilience through SIT

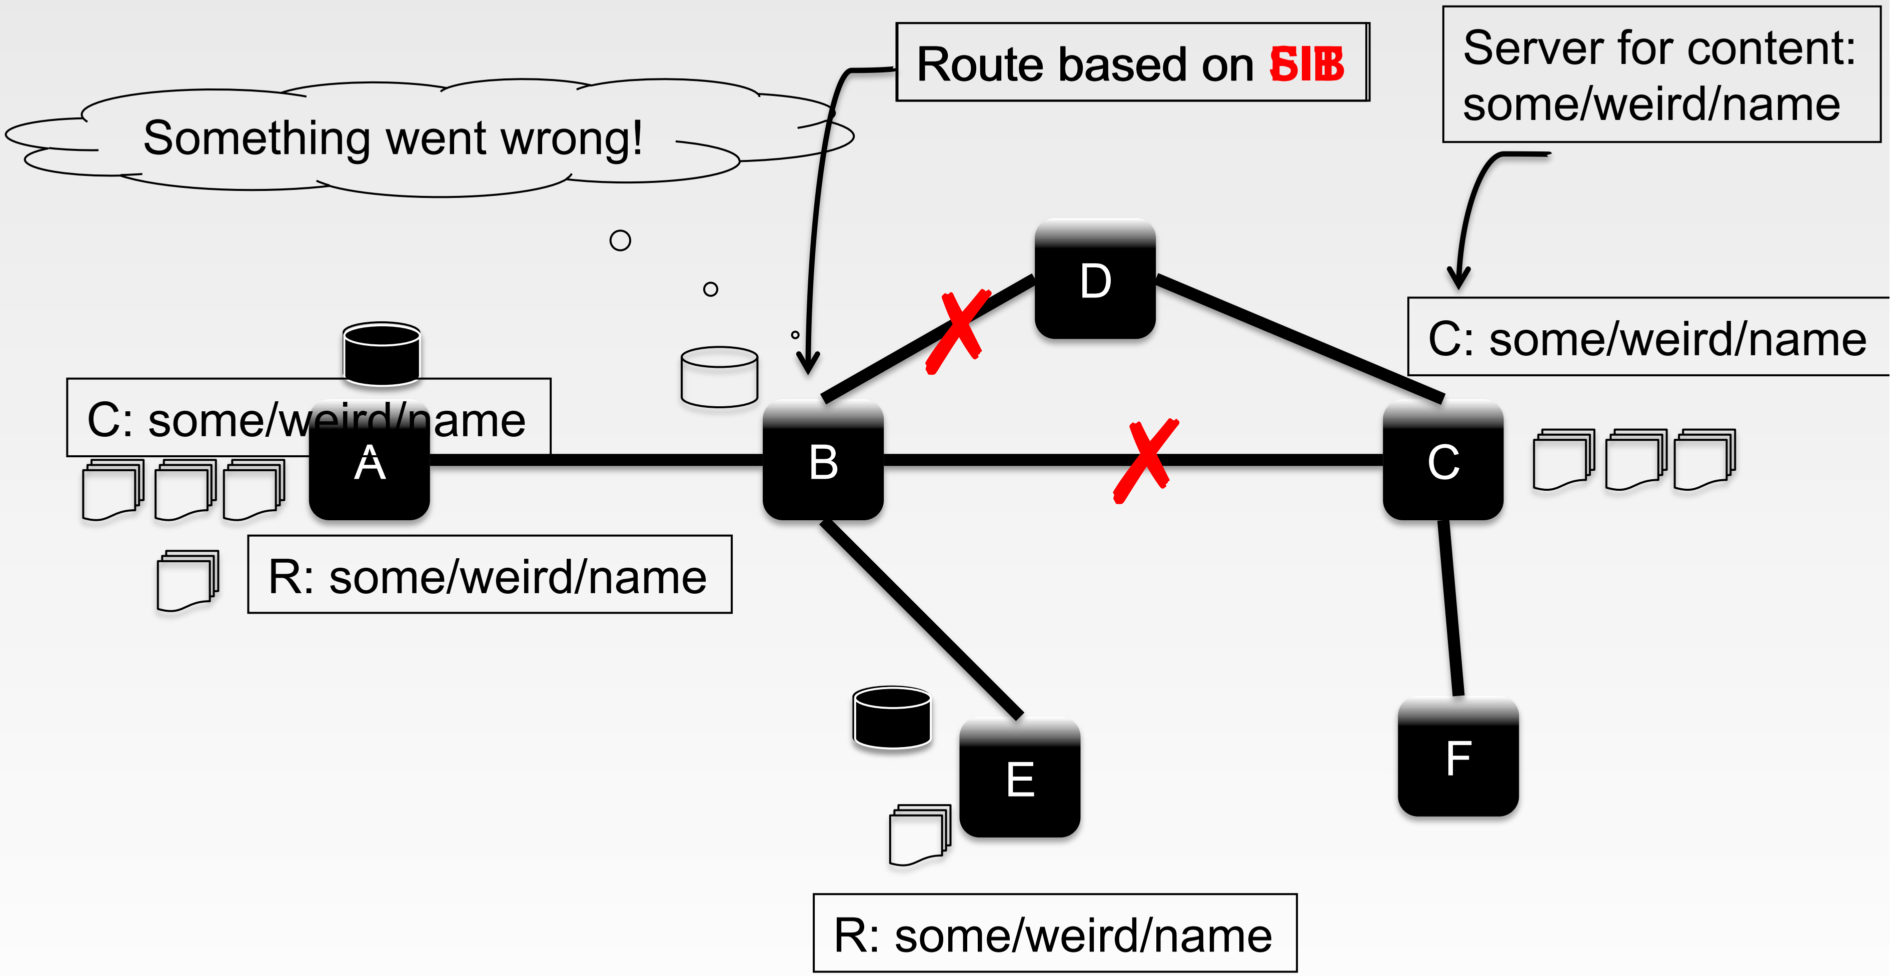

Key Design challenges & Contributions

- How to *augment the original NDN content router* to increase information resilience under fragmentation?
  - How to forward Interests when network fragments?
- What changes are required to the main *ICN packets format* and their processing in order to enable P2P-like content distribution?
- Can we *measure information resilience*?
  - We build Markov processes for the *hit probability and the time to absorption* of an item and find lower bounds

Router Design

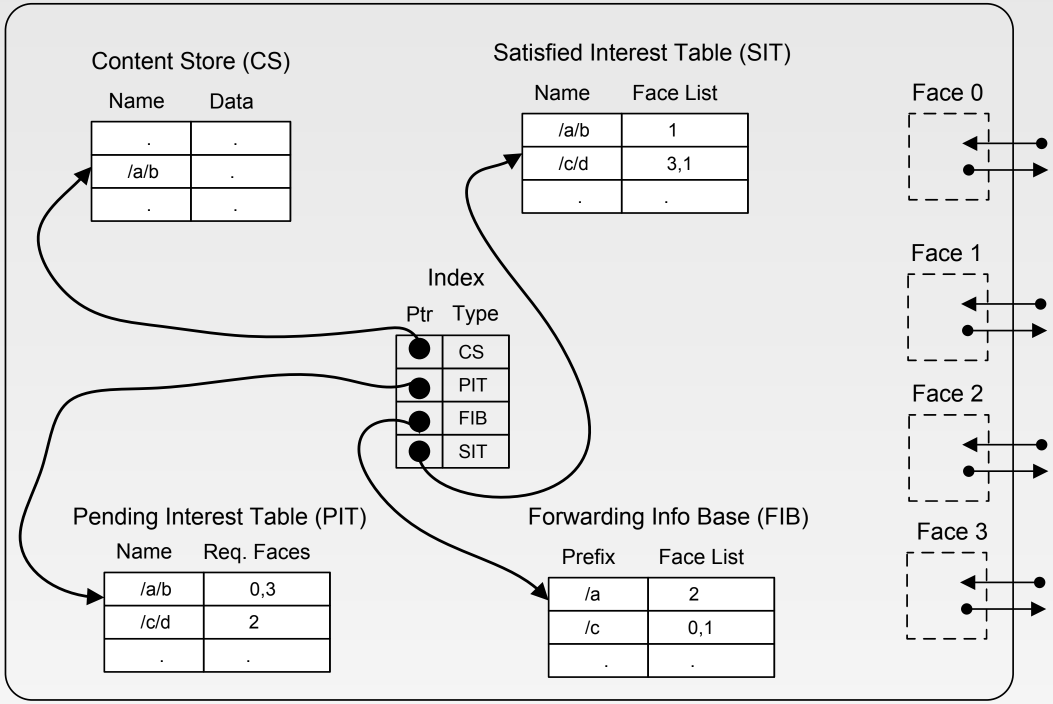

- Content Store (CS)
- Pending Interest Table (PIT)
- Forwarding Information Base (FIB)

Same to NDN original model

Satisfied Interest Table (SIT)

- Keeps track of data packet next hop.
- "*Breadcrumbs*" for user-assisted caching.
- Allows a list of outgoing faces.
- Similar to Persistent Interests (PI) in C. Tsilopoulos and G. Xylomenos, "Supporting Diverse Traffic Types in ICN" ACM SIGCOMM ICN 2011.

Packet Processing

- Interest Packet format
  - ***Destination flag (DF)*** bit to distinguish whether the Interest is headed towards content origin (DF=0), or towards neighbouring users (DF=1).
- Interest Packet processing
  - Normal operation (*i.e.,* no fragmentation): Same as in NDN
  - Fragmentation Detected: If the Interest cannot find a match in CS, PIT and FIB then *DF is set to 1 and follows entries in SIT*.
  - An Interest with DF=1 can be replied both by routers and by users with matching cached content.
- Data packet processing
  - Exactly the same as in NDN; follow the chain of PIT entries.
  - *A passing by Data packet installs SIT entries.*
  - Optionally cached in CS of each passing by router (under investigation).

Metrics

- **Satisfaction** *(% of issued interests*).
- **Absorbed Items** (*% of content items*).
- **Mean Absorption Time** (*sec*).
- **User Responses** (*% of satisfied interests*)
- **Minimum Hop Distance** (*hops*)
- **Traffic overhead** (*hops*)

Experiments

- Model validation
- Impact of cache size
- Impact of users’ disconnection rate.

Conclusions

- **Conceptual Gain:** A Downstream FIB can enable a *native content distribution network*
- **Performance Gain:**
  - A Downstream FIB can improve performance by reducing delay and load on core Internet links
  - Through a Downstream FIB, it is very easy to make the network resilient to fragmentation (at least in case of disasters). Popular content stays alive for many hours.
- **Implementation Considerations:** A Downstream FIB is not memory-intensive – acts like a cache.

Thanks!
Questions?

We'll soon have openings in our lab!

Ioannis Psaras
i.psaras@ucl.ac.uk
http://www.ee.ucl.ac.uk/~ipsaras/

Performance Bounds

System model

Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $|\mathcal{E}| = E$ and $|\mathcal{V}| = V$.

$\mathcal{M}$ uint size information items.

ζ_v and ϕ_v connection and disconnection rate of users at node v.

At each node v requests generated with rate $r_v = \{r_v^1, \ldots, r_v^M\}$,
r_v^m aggregate incoming request for item m.

ϑ_v^m Zipf law based popularity of item m at node v.

$$r_v^m = \zeta_v \cdot \vartheta_v^m = \zeta \cdot \vartheta_v^m = \zeta \cdot \frac{1/k^{z_v}}{\sum_{i=1}^{M} 1/i^{z_v}}$$

$\{X_m(t),\ 0 \le t < \infty\}$ the Markov process with stationary transition probabilities that depicts the number of users which have already retrieved item m and are connected in the network at time t.

$X_m(t)$ is a birth/death process with one absorbing state.

λ_n^m birth rate of the process at state n:

$$\lambda_n^m = \begin{cases} 0 & \text{if } n = 0, \\ \sum_{v \in \mathcal{V}} r_v^m = \sum_{v \in \mathcal{V}} \zeta \cdot \vartheta_v^m & \text{if } n > 0, \end{cases}$$

$\mu_n^m = n \cdot \phi_v = n \cdot \phi$; the death rate of the process.

Absorbing State Probability

$$u_s^m = \begin{cases} 1 & \text{if } \displaystyle\sum_{i=1}^{\infty} \rho_i^m = \infty, \\ \dfrac{\sum_{i=s}^{\infty} \rho_i^m}{1 + \sum_{i=1}^{\infty} \rho_i^m} & \text{if } \displaystyle\sum_{i=1}^{\infty} \rho_i^m < \infty. \end{cases}$$

where

$$\rho_i^m = \begin{cases} 1 & \text{if } i = 0, \\ \dfrac{\mu_1^m \mu_2^m \cdots \mu_i^m}{\lambda_1^m \lambda_2^m \cdots \lambda_i^m} = \dfrac{\phi \cdot 2\phi \cdots i\phi}{\lambda^m \lambda^m \cdots \lambda^m} = \left(\dfrac{\phi}{\lambda^m}\right)^i \cdot i! & \text{if } i > 0. \end{cases}$$

Mean Time to Absorption

$$T_s^m = \begin{cases} \infty & \text{if } \sum_{i=1}^{\infty} \frac{1}{\lambda_i^m \cdot \rho_i^m} = \infty, \\ \sum_{i=1}^{\infty} \frac{1}{\lambda_i^m \cdot \rho_i^m} + \sum_{k=1}^{s-1} \rho_k^m \sum_{j=k+1}^{\infty} \frac{1}{\lambda_j^m \cdot \rho_j^m} & \text{if } \sum_{i=1}^{\infty} \frac{1}{\lambda_i^m \cdot \rho_i^m} < \infty. \end{cases}$$

- **Result:** *When the death rate of the users interested in a content item is larger than the corresponding birth rate, the item will finally get absorbed when the content origin is not reachable.*
  - The formula above gives us the *"time to absorption"*

[1] H. M. Taylor and S. Karlin, "An Introduction to Stochastic Modeling, 3rd edition", Academic Press, 1998.

Performance Evaluation

Strategies/Policies (after the network fragmentation)

- **Interest forwarding policies**
  - SIT based forwarding policy (STB)
  - Flooding forwarding policy (FLD)
- **Caching policies**
  - No caching policy (NCP)
  - Edge caching policy (EDG)
  - En-route caching policy (NRT/LCE)
- **Placement/Replacement policies**
  - Least Recently Used policy (LRU)

Evaluation setup

- *Tool:* Icarus
- *Network topology:* 50 nodes - Internet topology Zoo
- *Traffic demand:* 1*req/sec* at each node
- *Request distribution:* Zipf *and* localised, *i.e.,* different across different regions
- *Connection rate:* 1 new user per sec
- "Initialization period" of 1 *hour.* "Observation period" of 3 *hours.* Network fragmentation and origin servers of all items are not reachable.

Model Validation

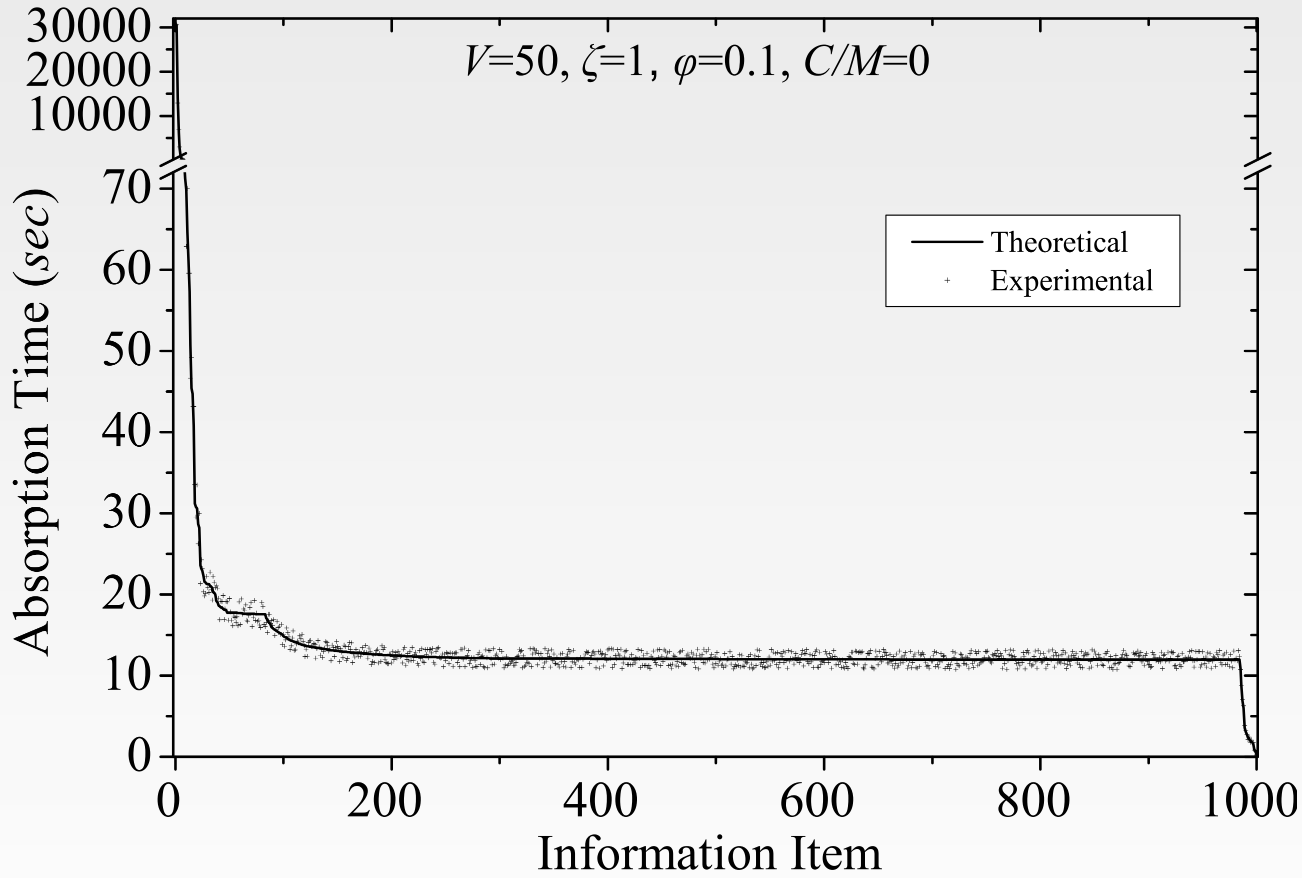

Perfect match between model and simulation!

Impact of the cache size

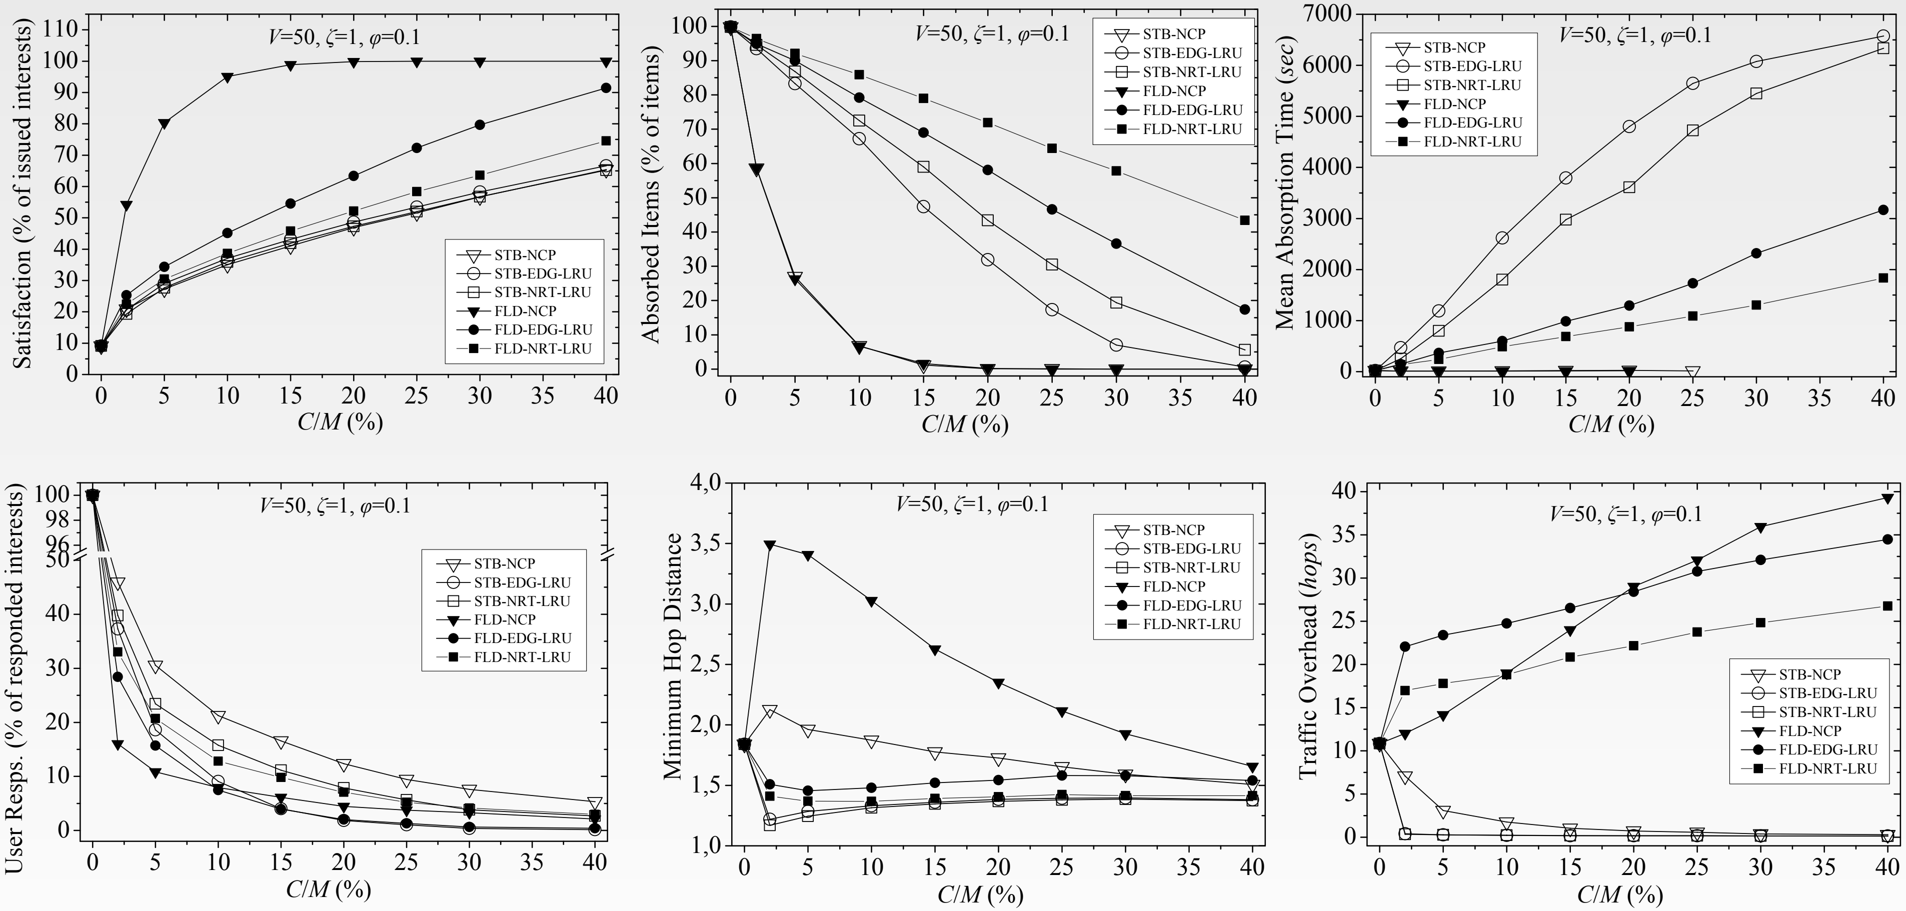

Popular messages can stay in the network for hours even with modest amounts of cache.

UCL

Impact of users' disconnection rate

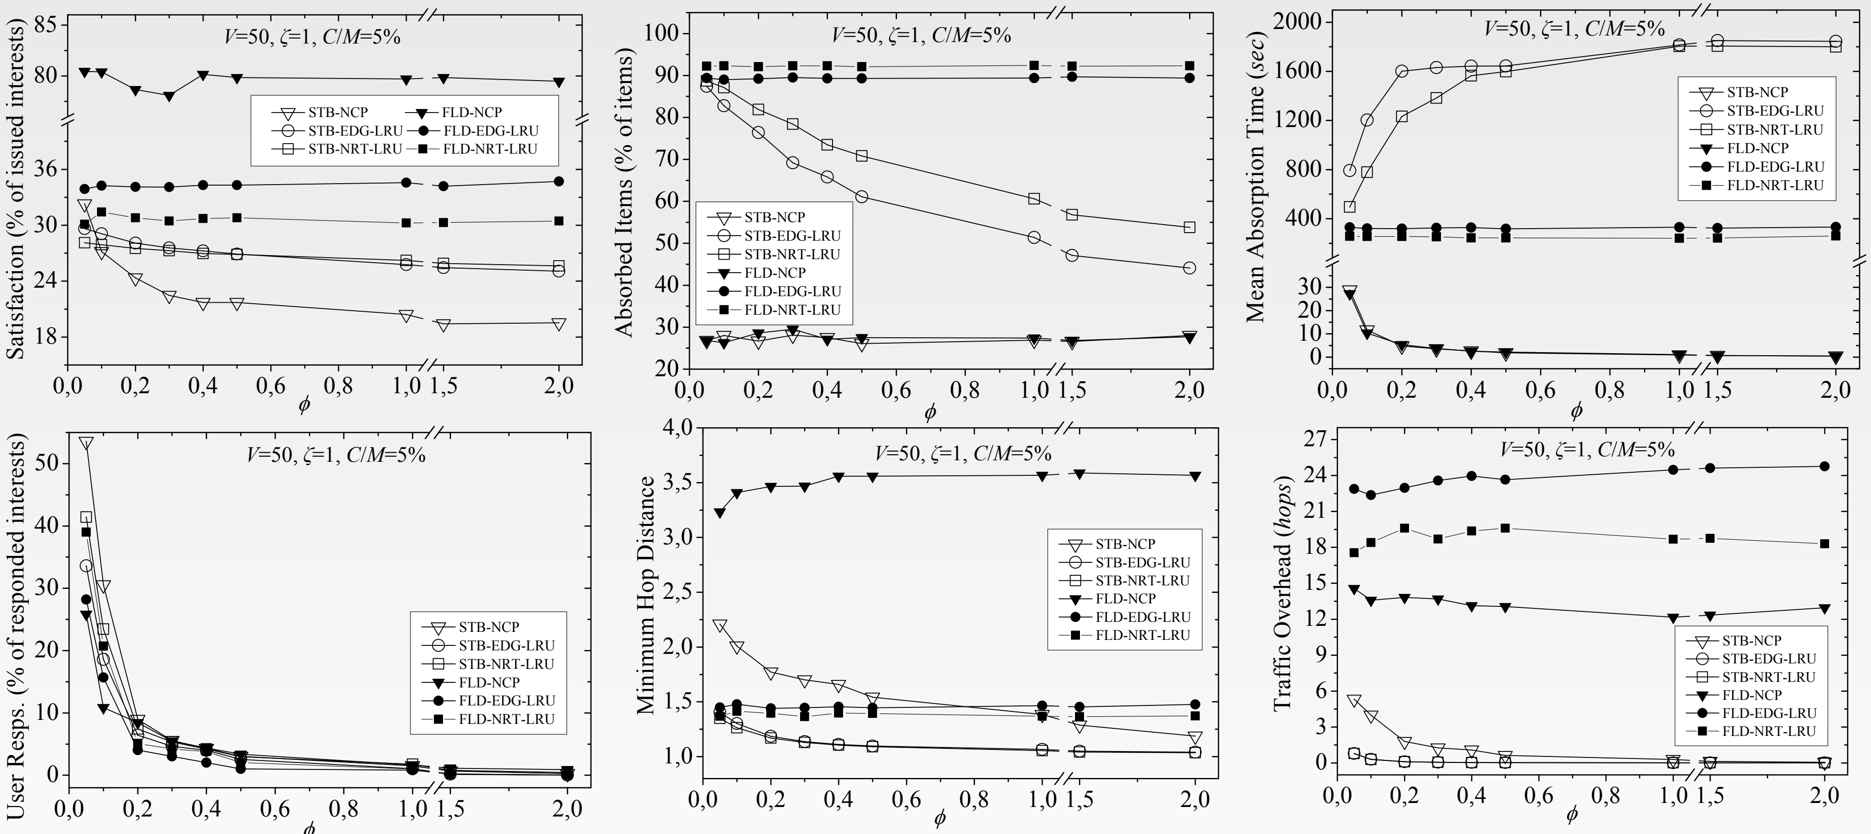

- When disconnection rate is larger than 0.2, less than 5% of the satisfied interests are served from users.
- The STB enabled mechanisms discard less popular items fast and maintain the rest items for a longer period.