Coupling Caching and Forwarding: Benefits, Analysis & Implementation

http://www.enst.fr/~drossi/ccnSim

Dario Rossi
dario.rossi@enst.fr

Giuseppe Rossini
giuseppe.rossini@enst.fr

Context

- Information centric neworks
  - Mobility, naming, security, …, and Caching
- Caching on Information centric networks
  - Not a goal on its own, but a means (e.g., congestion-aware to reduce load on bottleneck links, or cost-aware to reduce ISPs opex)
  - General viewpoint: Efficiently allocate/use caching resources
- Important notions
  - Temporary vs Persistent
  - On-path vs Off-path

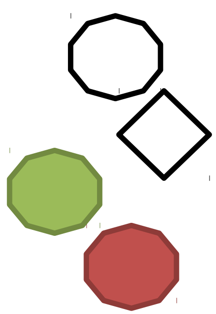

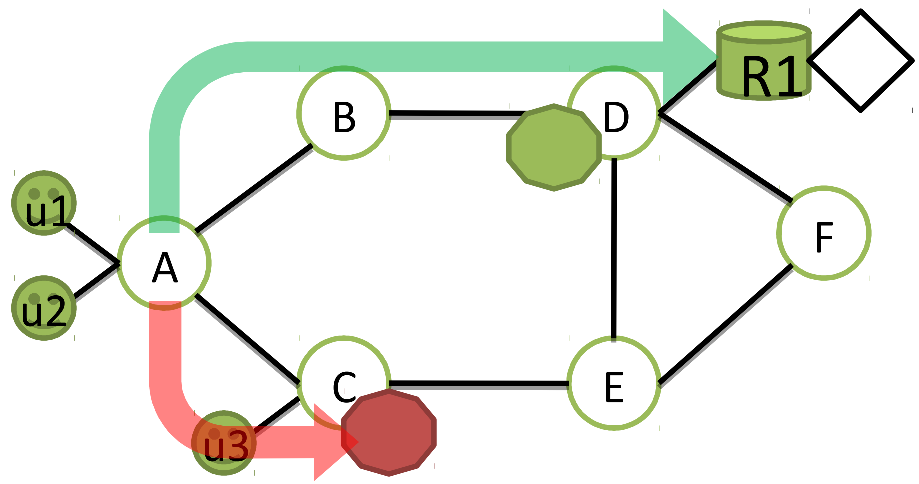

- This talk
  - How to efficiently find the closest,possibly off-path, content replica

Terminology

ICN design space is wide

- We may use FIB toward custodian if available
- We ignore proactive/periodic dissemination of cached copies availability on control plane
- We focus on reactive discovery of cached copies on the data plane

Algorithmically speaking, a cache network can be identified as a triple $\langle \mathcal{F}, \mathcal{D}, \mathcal{R} \rangle$

- ***F*** is the Forwarding strategy
- ***D*** is the cache Decision policy
- ***R*** is the cache Replacement policy

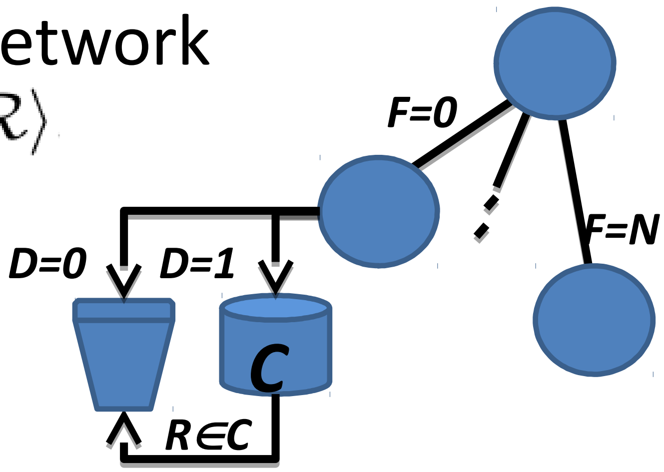

Related work

- Typically, focus on either *F* or *D*, but not both (*R*=LRU as replacement policy makes sense)

Meta-caching $\mathcal{D}$	Type	Knob	Ref
LCE			[6–13, 15–17, 23–25, 29–31, 35, 39]
Fix	Prob.	p	[5, 25, 35]
ProbCache	Prob.	Distance	[29]
LCD	Det.	Distance	[24, 25, 35]
WAVE	Prob.	Distance	[13]
Btw	Det.	Centrality	[9]

F=SPR by default

Forwarding $\mathcal{F}$	Type	Knob	Ref
SPR			
Source routing			
Flooding			
INFORM			
CATT			
NDN			
iNRR	[illegible]		[illegible]

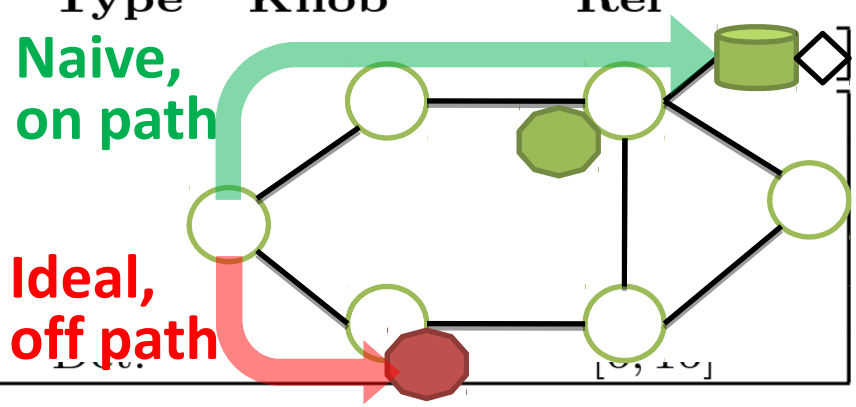

D=LCE by default

Agenda

- Benefits
  - Interest of joint <F,D>
  - Comparison with several CDN strategies [16]

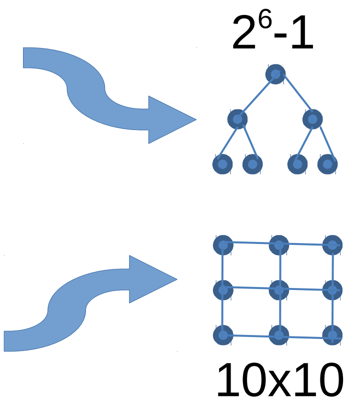

- Analysis
  - Modeling of <iNRR,LCE,LRU>
  - Extend <SPR,LCE,LRU>, aka aNET [31]
- Implementation
  - Notion of “Meta-interest” to avoid pollution across paths
  - Spoiler: Arbitrarely close to iNRR, slight delay tradeoff

[16] S. K. Fayazbakhsh, Y. Lin, A. Tootoonchian, A. Ghodsi,T. Koponen, B. M. Maggs, K. Ng, V. Sekar, and S. Shenker. Less pain, most of the gain: Incrementally deployable. In ACM SIGCOMM, 2013.

[31] E. J. Rosensweig, J. Kurose, and D. Towsley. Approximate Models for General Cache Networks. In IEEE INFOCOM, 2010

Coupling Caching and Forwarding: **Benefits,** Analysis & Implementation

(scenario details in the paper)
(scenario scripts available at)
http://www.enst.fr/~drossi/ccnSim

Part I

CDN vs ICN scenario

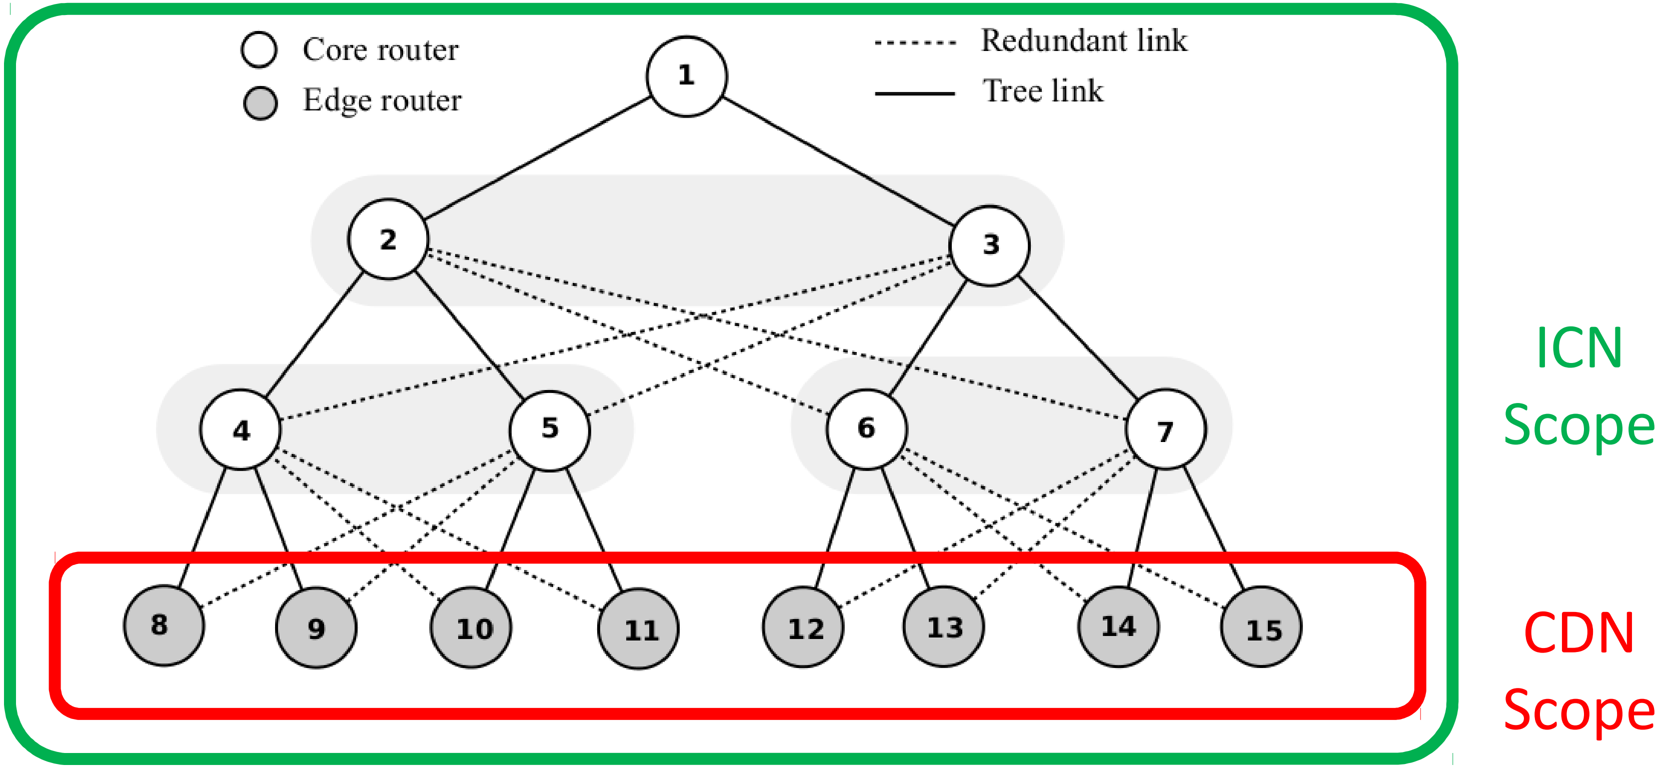

Fig. 1. Redundant 4-level binary tree. Dashed links are present with probability μ. Shadowed blocks represent aggregate caches seen by lower level nodes in presence of redundancy ($\mu > 0$).

CDN Policies

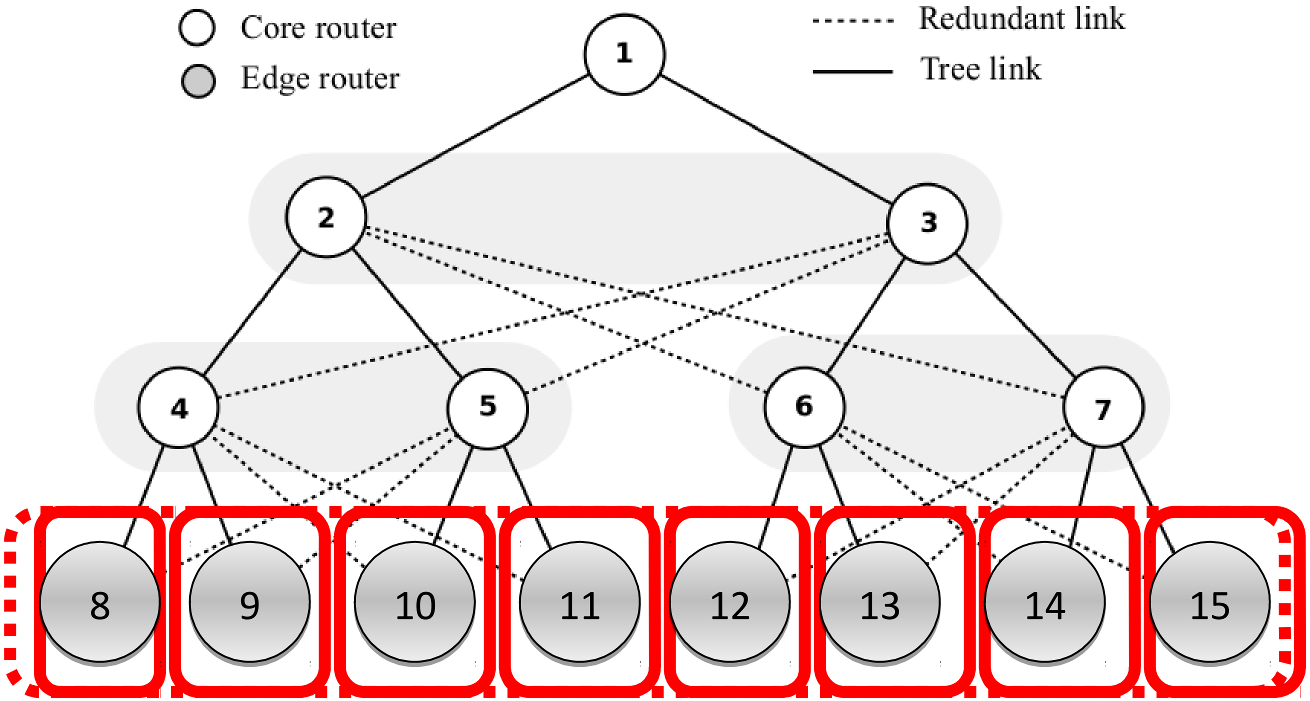

Edge Naive CDN, with all caches operating independently

EdgeCoop Smarter CDN, all caches cooperate (with iNRR)

EdgeNormCoop As EdgeCoop, but individual caches have double size (network-wide same budget w.r.t. ICN scope)

ICN Policies

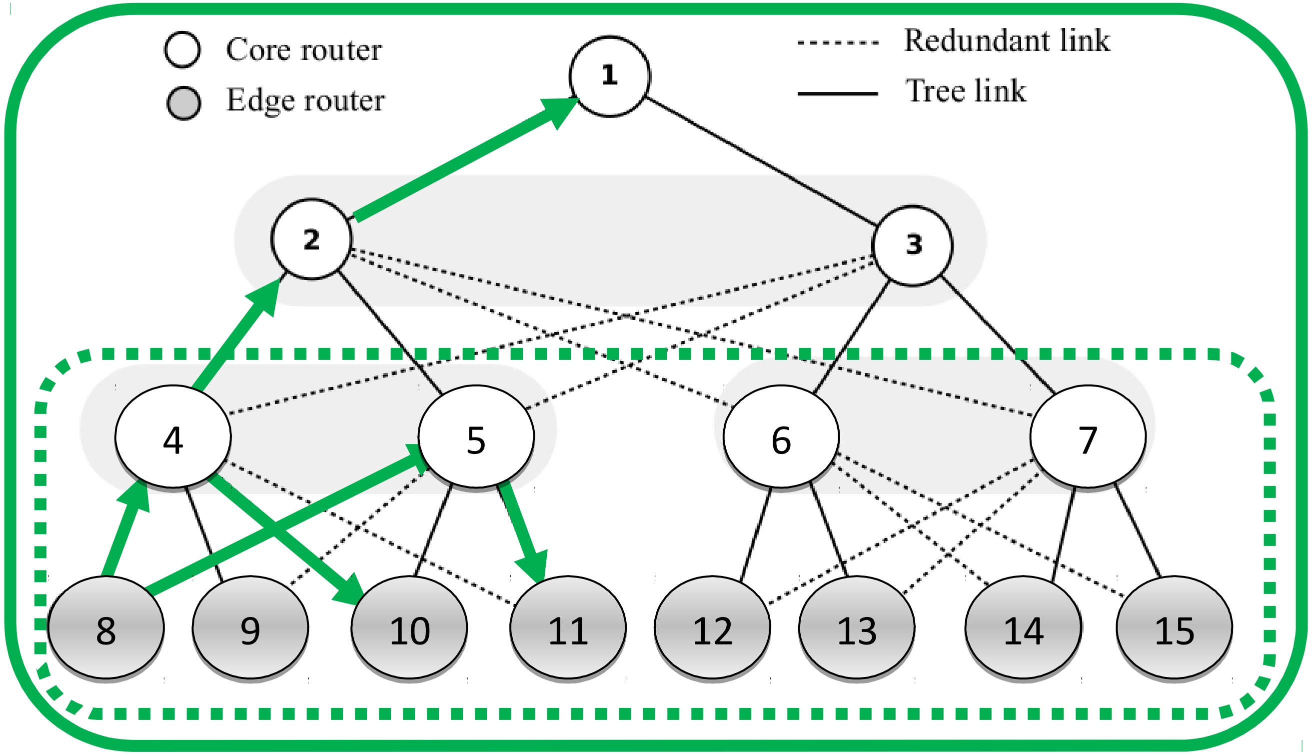

<SPR,LCE> Naive ICN: on-path caching toward repository

<NRR,LCE> Access to off-path copies, but cache pollution due to LCE

<NRR,LCD> Smarter ICN: off-path copies, limits cache pollution via LCD

Configurations: 2-levels vs Ubiquitous caches
(as for EdgeNormCoop, keeping the same overall cache budget)

CDN vs ICN policies

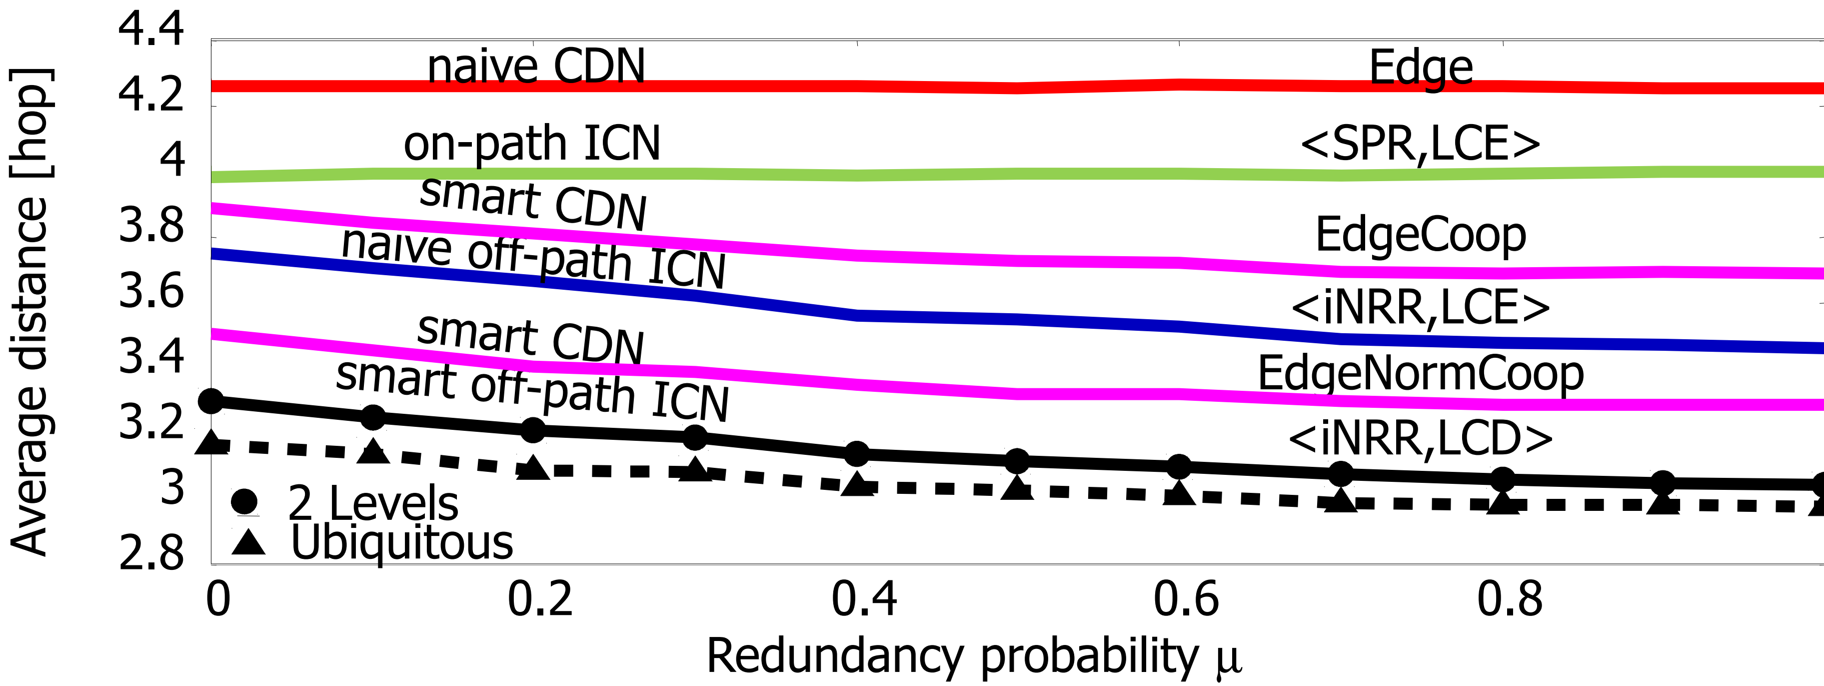

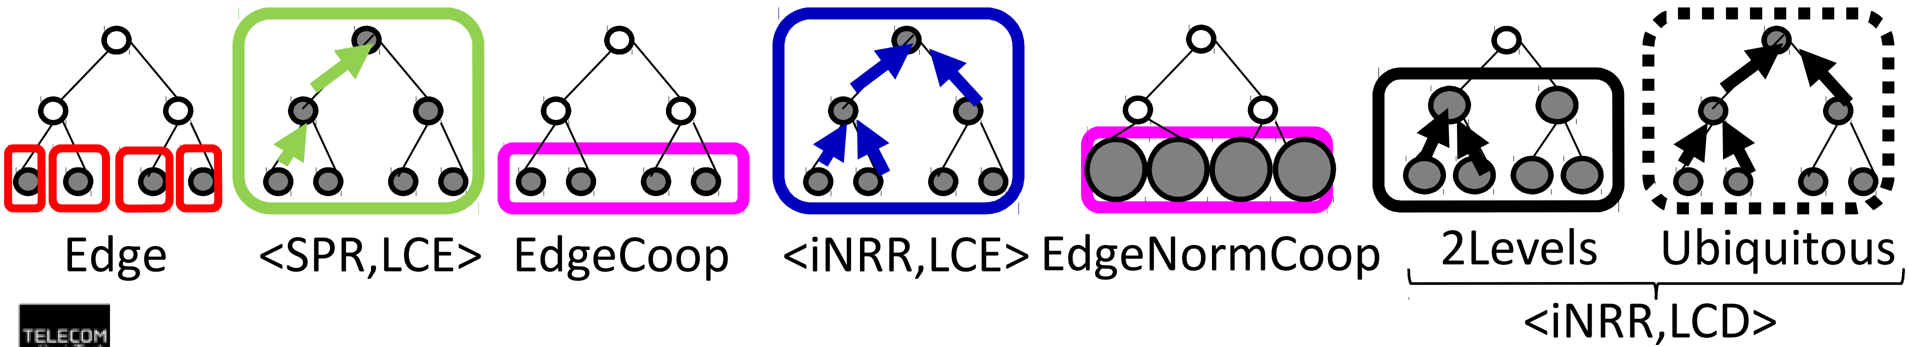

CDN vs ICN policies

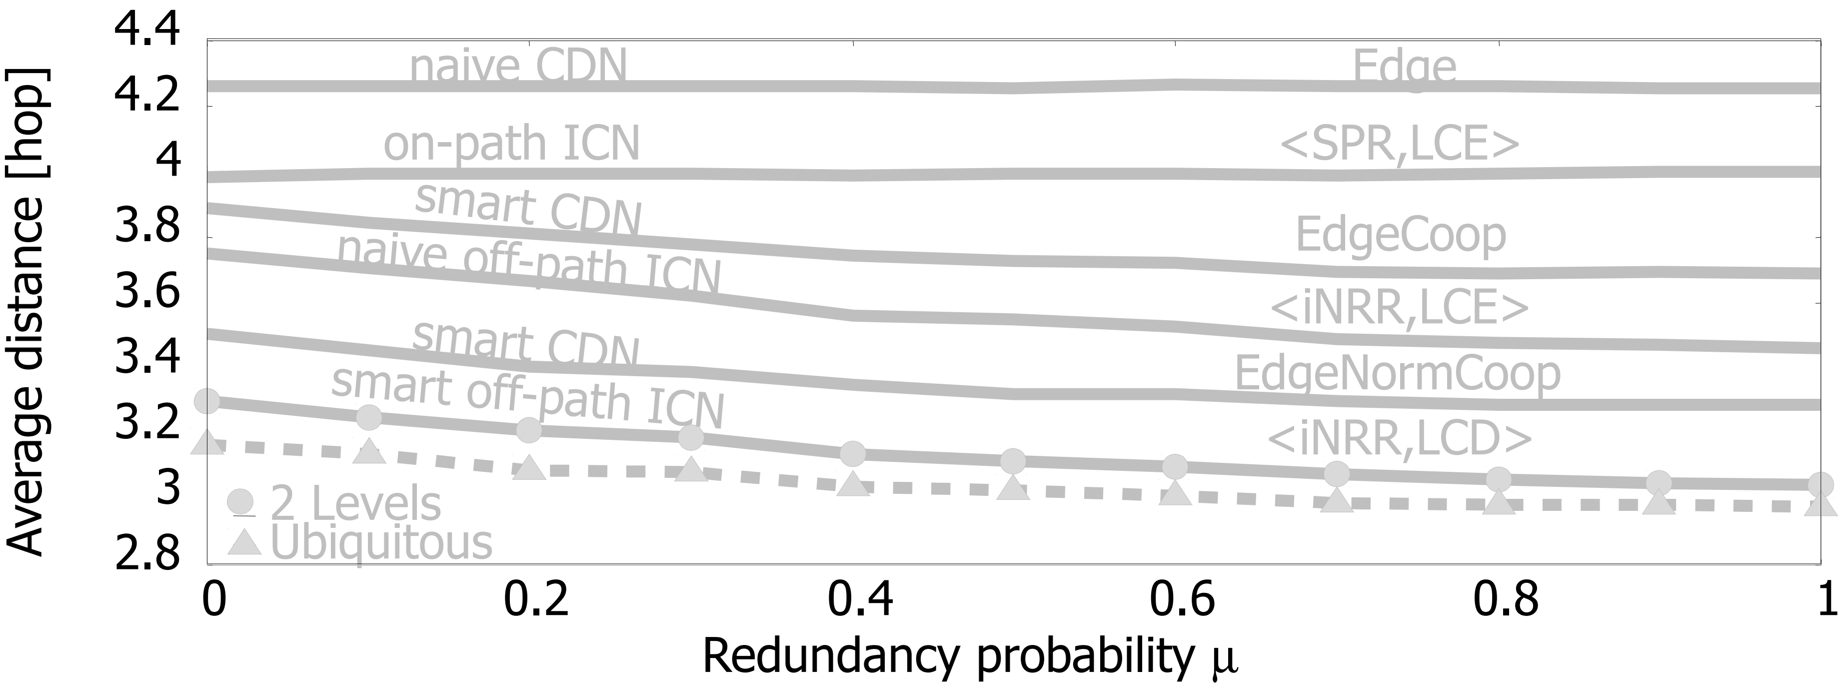

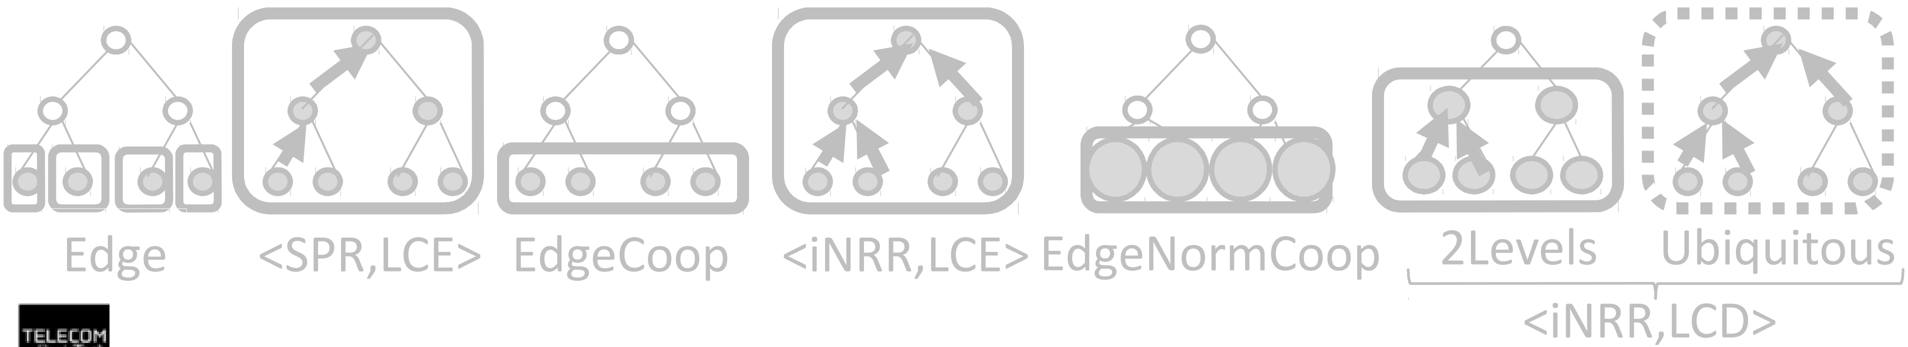

CDN vs ICN policies

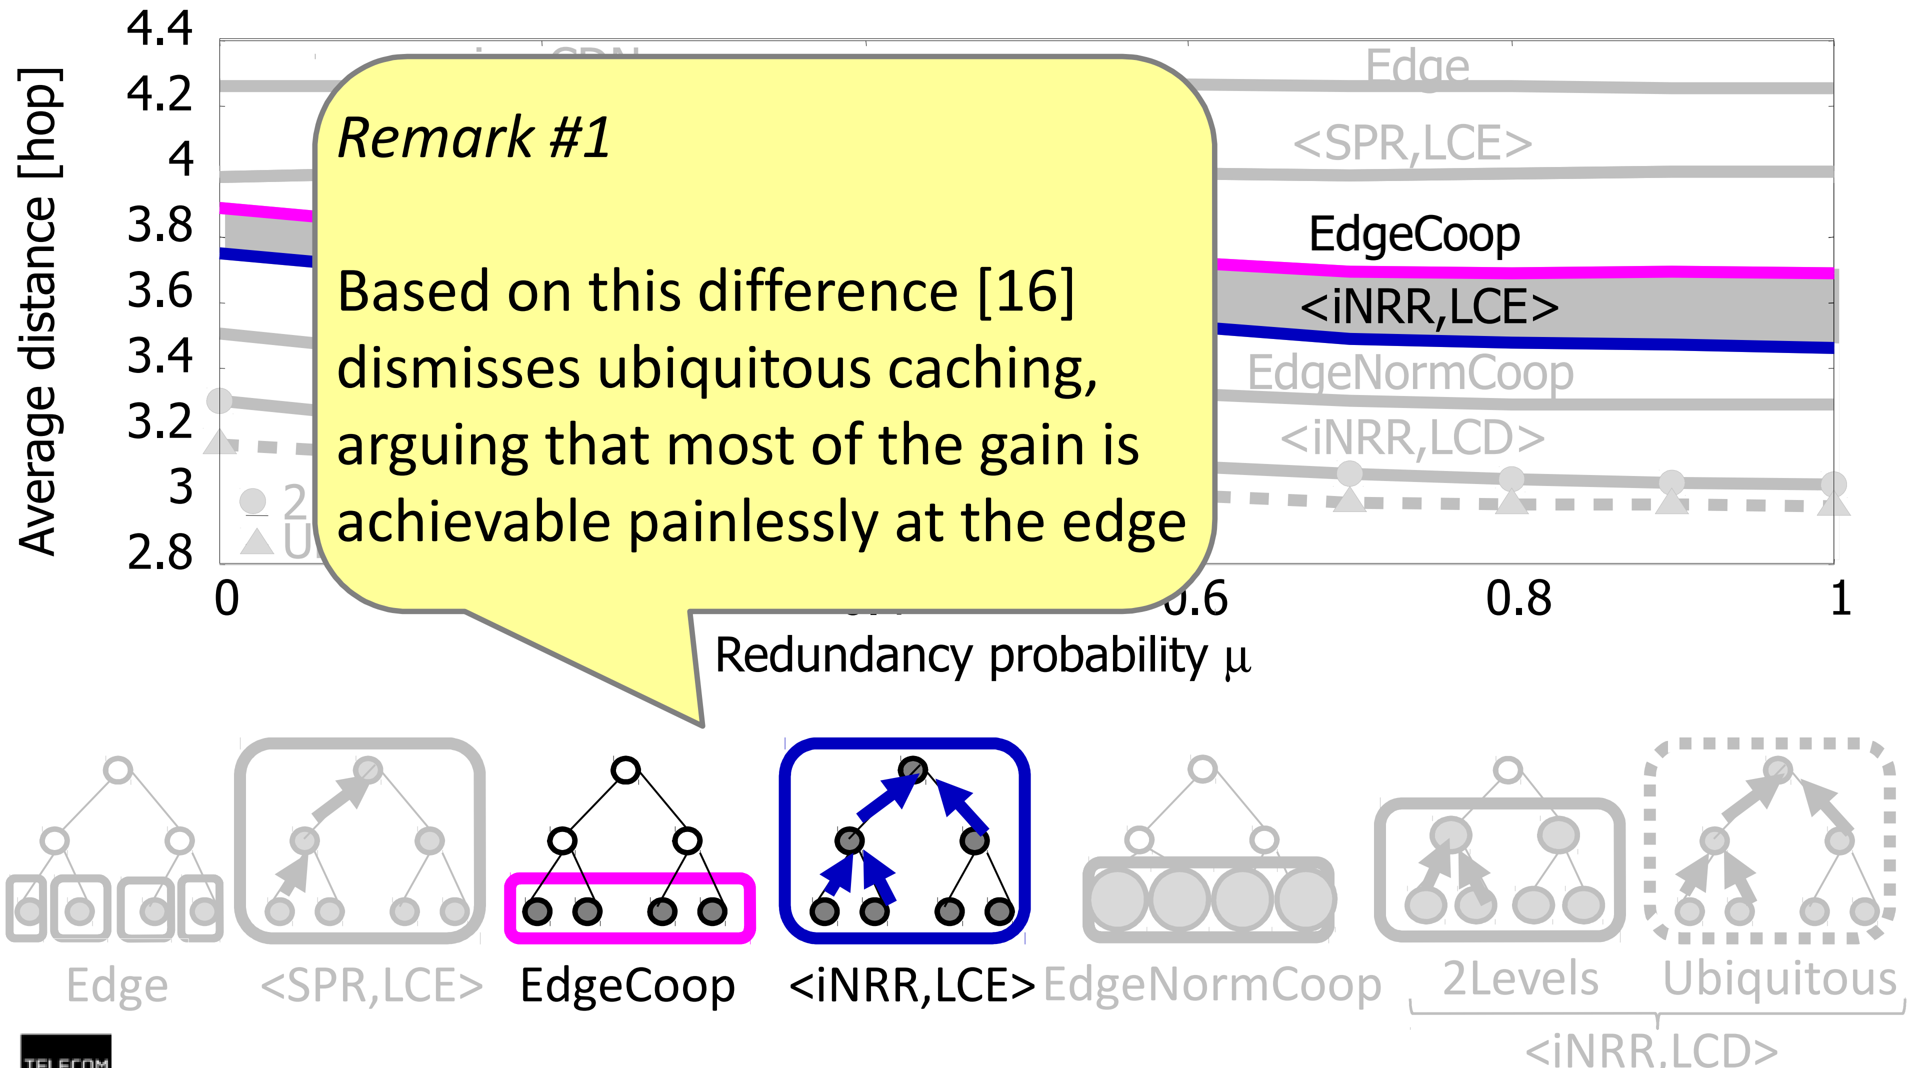

Remark #1

Based on this difference [16] dismisses ubiquitous caching, arguing that most of the gain is achievable painlessly at the edge

CDN vs ICN policies

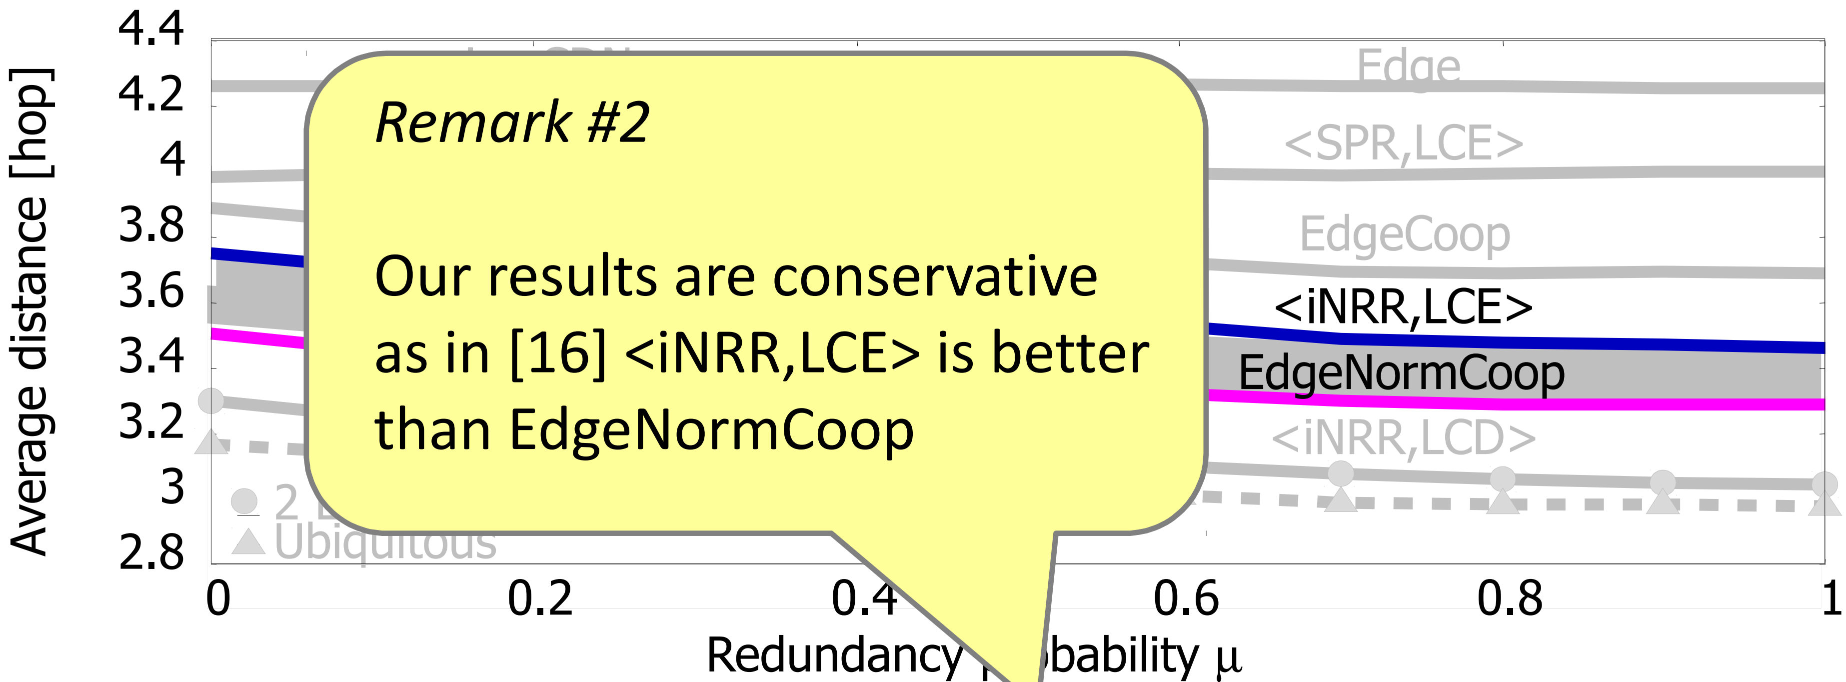

Remark #2

Our results are conservative as in [16] <iNRR,LCE> is better than EdgeNormCoop

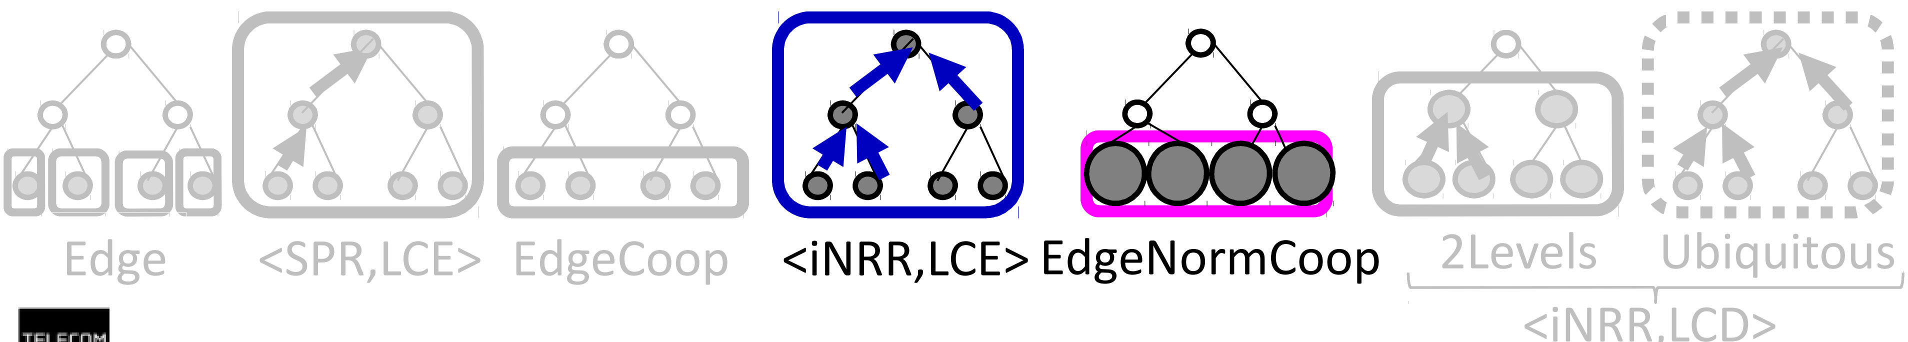

CDN vs ICN policies

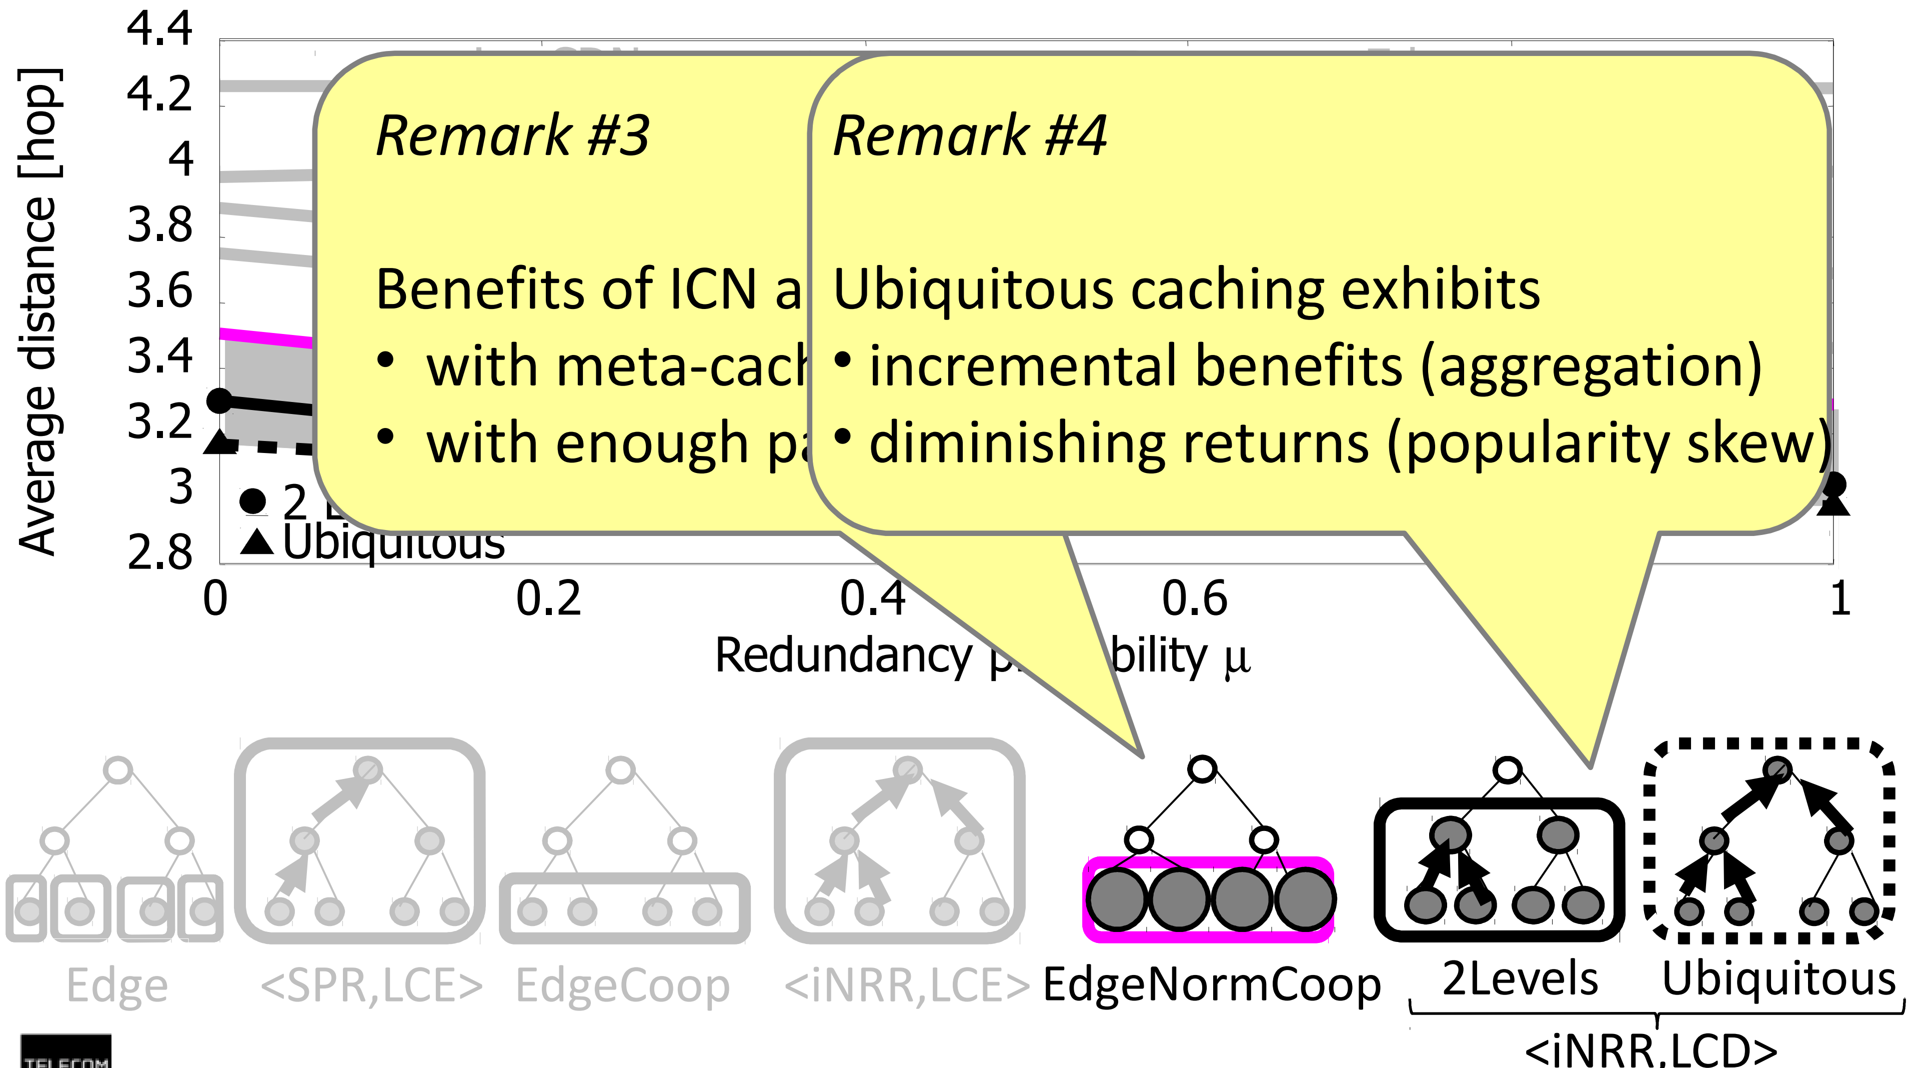

CDN vs ICN policies

Remark #5

- Technico economic studies needed to assess if benefits justify deployment
- ICN success unlikely driven *only* by caching benefits

μ	⟨SPR, LCE⟩	Edge Coop	⟨iNRR, LCE⟩	Edge Norm Coop	⟨iNRR, LCD⟩ 2-Levels	⟨iNRR, LCD⟩ Ubiquitous

Redundancy probability μ

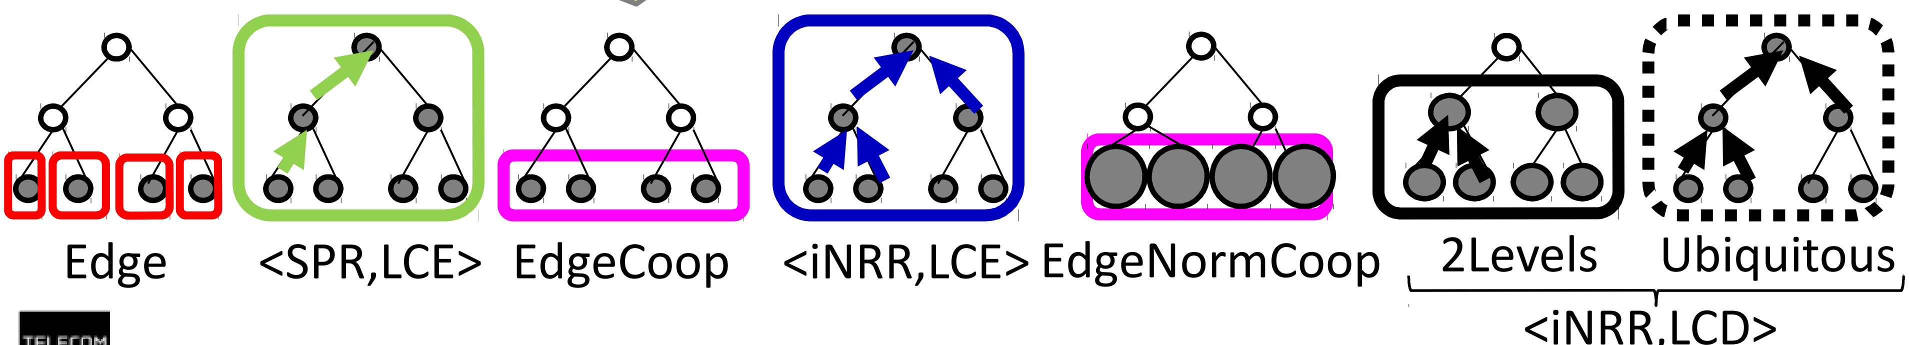

Coupling Caching and Forwarding: Benefits, **Analysis** & Implementation

Part II

Modeling iNRR

- The Ideal Nearest Replica Routing (iNRR) [2] is nice and worth *modeling*
- Let start from the aNET model [31] for SPR

$$r_{i,v} = \lambda_{i,v} + \sum_{u:R(u,\mathcal{S}(i))=v} m_{i,u} \quad (1)$$

$$p_{i,v} = \frac{r_{i,v}}{\sum_{j=1}^{N} r_{jv}} \quad (2)$$

$$\vec{\pi}_v = LRU(\vec{p}_v, |v|) \quad (3)$$

$$m_{i,v} = r_{i,v}(1 - \pi_{i,v}) \quad (4)$$

Modeling iNRR

$$r_{i,v} = \lambda_{i,v} + \sum_{u:u \in N(v)} m_{i,u,v} \quad (5)$$

$$p_{i,v} = \frac{r_{i,v}}{\sum_{j=1}^{N} r_{jv}} \quad (6)$$

$$\vec{\pi}_v = LRU(\vec{p}_v, |v|) \quad (7)$$

$$m_{i,v} = r_{i,v}(1 - \pi_{i,v}) \quad (8)$$

Similar to aNET

$$s_{i,v,u} = \sum_{\substack{x:R(v,x)=u \\ \wedge x \in B(v,\mathcal{S})}} \left[\prod_{y \in B_i(v,x)} (1 - \pi_{i,y}) \right] \frac{\pi_{i,x}^2}{\sum_{z \in B_b(v,x)} \pi_{i,z}} \quad (9)$$

Define split ratio among multiple path

$$m_{i,v,u} = \begin{cases} m_{i,v} s_{i,v,u} & u \neq R(v,\mathcal{S}) \\ m_{i,v}(1 - \sum_{w \neq u} s_{i,v,w}) & u = R(v,\mathcal{S}) \end{cases} \quad (10)$$

Applies split ratio to miss stream

Model accuracy

aNET suffers from IRM violation due to miss-stream far from the leafs

LRU approximation very accurate for leaf nodes

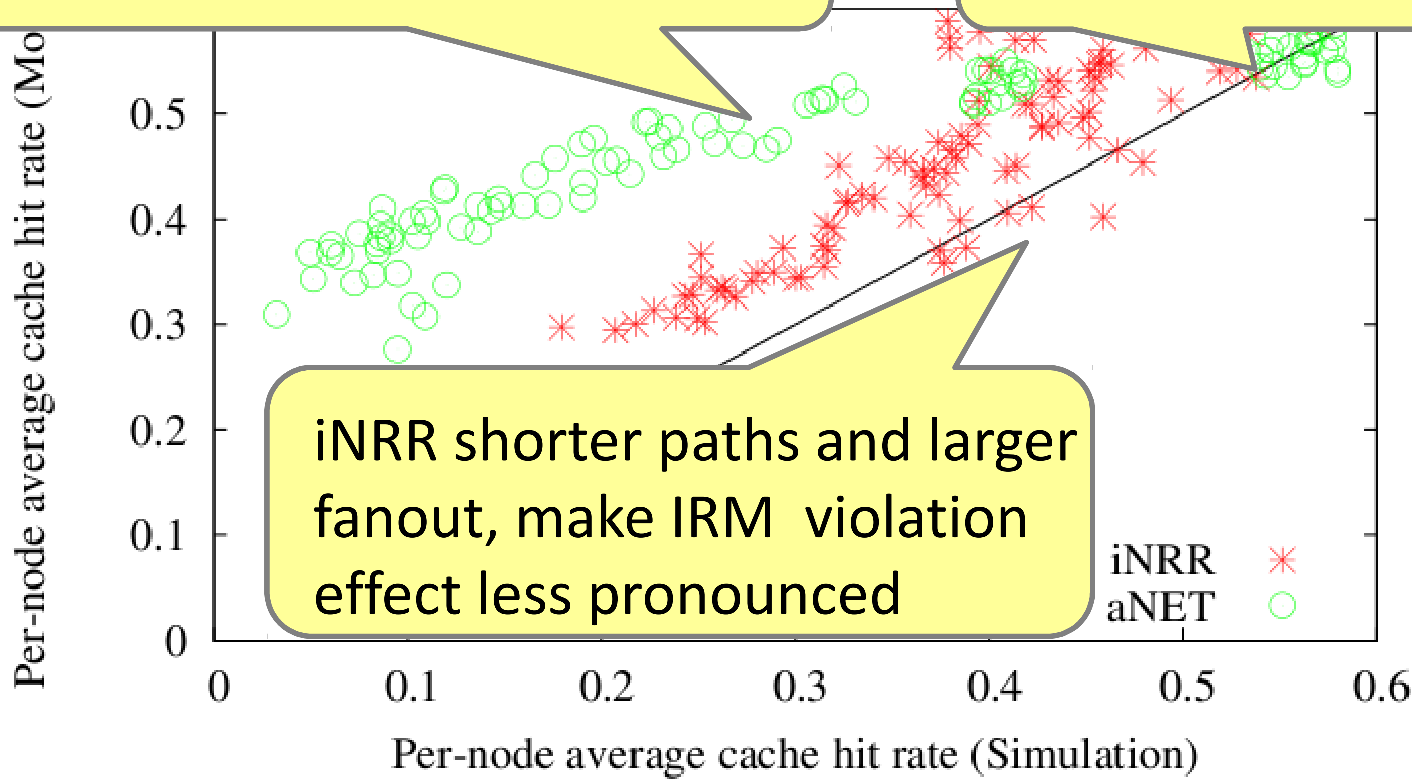

iNRR shorter paths and larger fanout, make IRM violation effect less pronounced

Fig. 7. Scatter plot of the average cache hit per node $\bar{\pi}_v$ obtained via simulation vs model, for aNET and iNRR, on a 10x10 grid.

Coupling Caching and Forwarding: Benefits, Analysis & **Implementation**

Part III

Architecture

- General scheme: forward interests along either
  - FIB for persistent on-path copies, proactively advertised by routing
  - TFIB for temporary off-path cached copies, reactively discovered via exploration

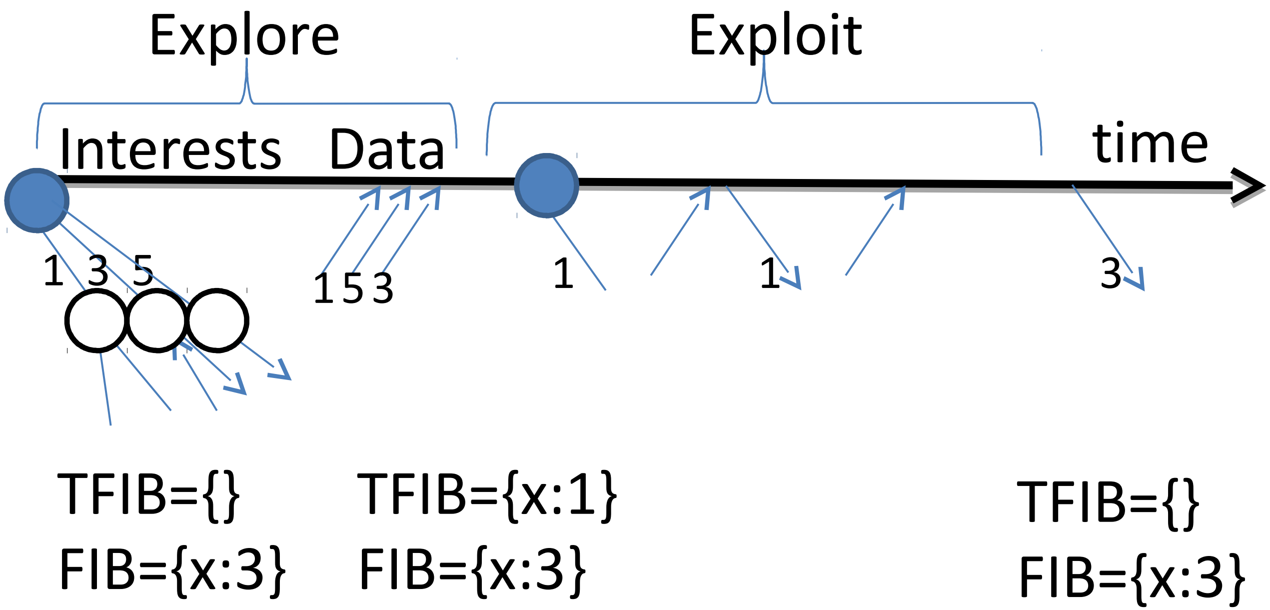

Meta-interests

- The Ideal Nearest Replica Routing (iNRR) [2] is nice and worth *implementing*
- Explore via scoped flooding, at distance d hops

1-phase design (NRR')

Use regular interest packets

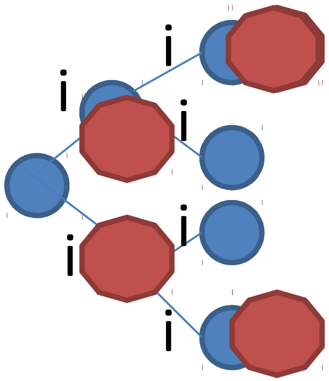

All cache hits return real data.
Redundant for popular content,
Useless for unpopular content

2-phase design (NRR'')

Use meta-interest packets

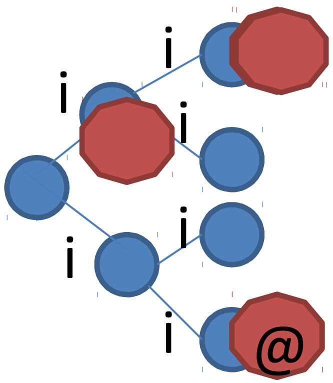

Hits indicate data @vailability.
Afterward, actual requests only sent through a single TFIB path

Practical NRR very close to iNRR

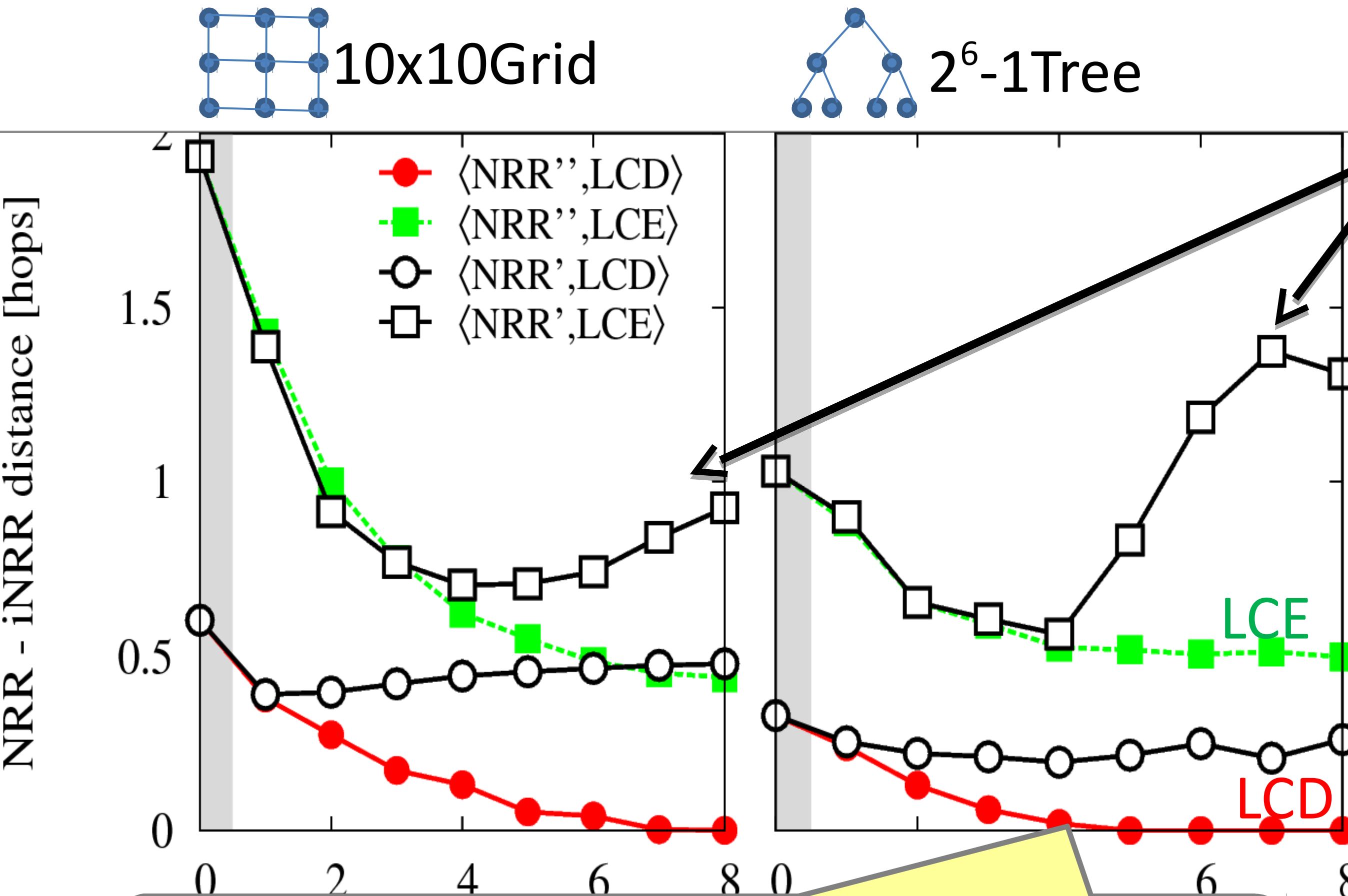

1-phase
Regular interests yield to cache pollution penalty

2-phases
Meta-interests confine pollution to a single path

Meta-caching limit pollution along the path

Meta-interests + LCD meta-caching = <NRR'',LCD> arbitrarily close to iNRR

Implementing NRR

- Implementing ideal Nearest Replica Routing (iNRR)
- Explore via scoped flooding, max distance *d* hops

1-phase design (NRR')

Use regular interest packets

Regular interest generate cache decision, triggering cache replacement and pollution

Fastest propagation, time to first chunk is lowest

2-phase design (NRR'')

Use meta-interest packets

Meta-interest return only data-availability indication, not data itself. Avoid Pollution!

Slower propagation, time to first chunk longer (subsequent chunks are found at NRR)

Implementing NRR

- Implementing ideal Nearest Replica Routing (iNRR)
- Explore via scoped flooding, max distance *d* hops

In practice, delay penalty

- only affects the first chunk, upper-bounded by network RTT

Delay not necessarily longer

- as content is closer, 1st chunk delay may even be shorter

replacement and pollution

not a... ...oid Pollution!

Fastest propagation, time to first chunk is lowest

Slower propagation, time to first chunk longer (subsequent chunks are found at NRR)

Coupling Caching and Forwarding: Benefits, Analysis & Implementation

Aftermath

Conclusions

- Contributions
  - Benefits, analysis and implementation of joint forwarding and caching
  - *Meta-caching*:
    implicit coordination of caching decisions *along a single path*
  - *Meta-interests*:
    implicit coordination of caching decisions *across multiple paths*
- Limits
  - Can we do better (e.g., 2-LRU instead of LCD) ?
  - How far are is <iNRR,LCD,LRU> from optimum ?
  - Only LCE is modeled
  - Only synthetic topologies are simulated, with
    stationary popularity and no spatial skew

Implications

- Comparison guidelines
  - Algorithmic suggestion for IRTF ICNRG baseline comparison
  - Good practice = include <iNRR,LCD,LRU> (or better)
  - Bad practice = using only naive <SPR,LCE,LRU>

- Comparison is feasible
  - NRR'' implementation available in existing tools (e.g., ccnSim)
  - iNRR (and NRR'') simple to implement in other simulators
  - Results consistent *across* simulators: Come see the demo!

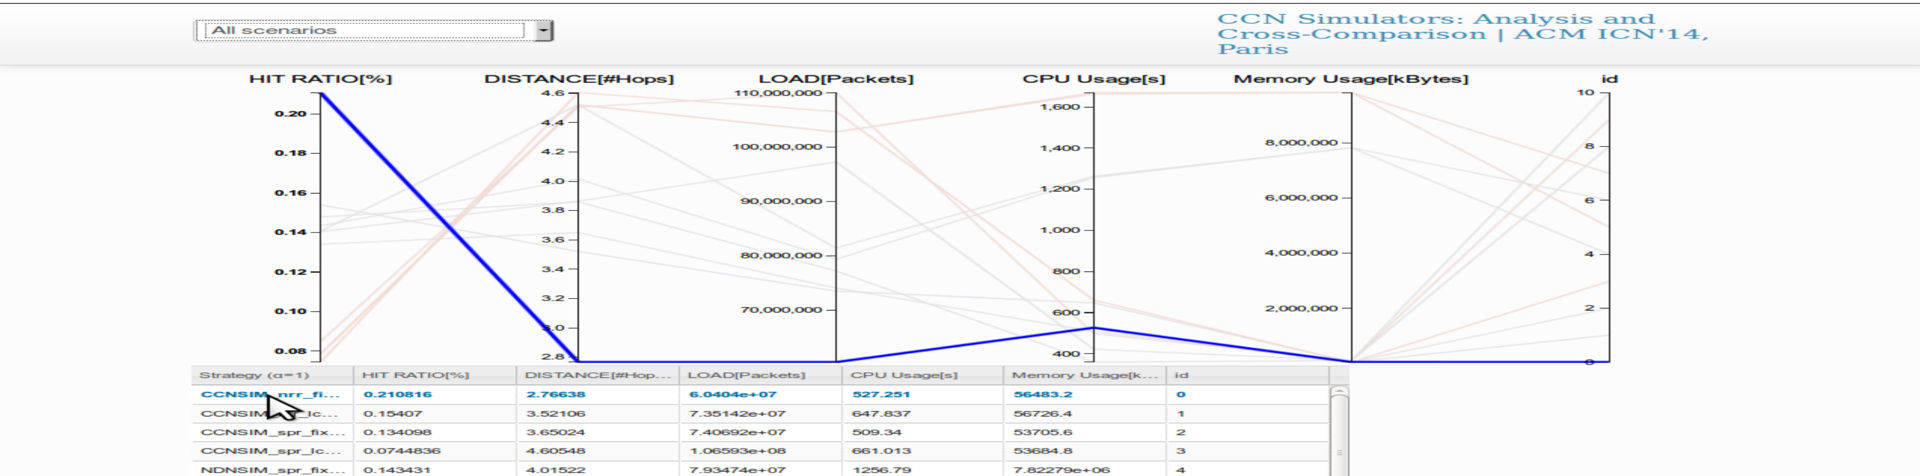

?? || //

Backup slides

Interest of joint <F,D,(R=LRU)>

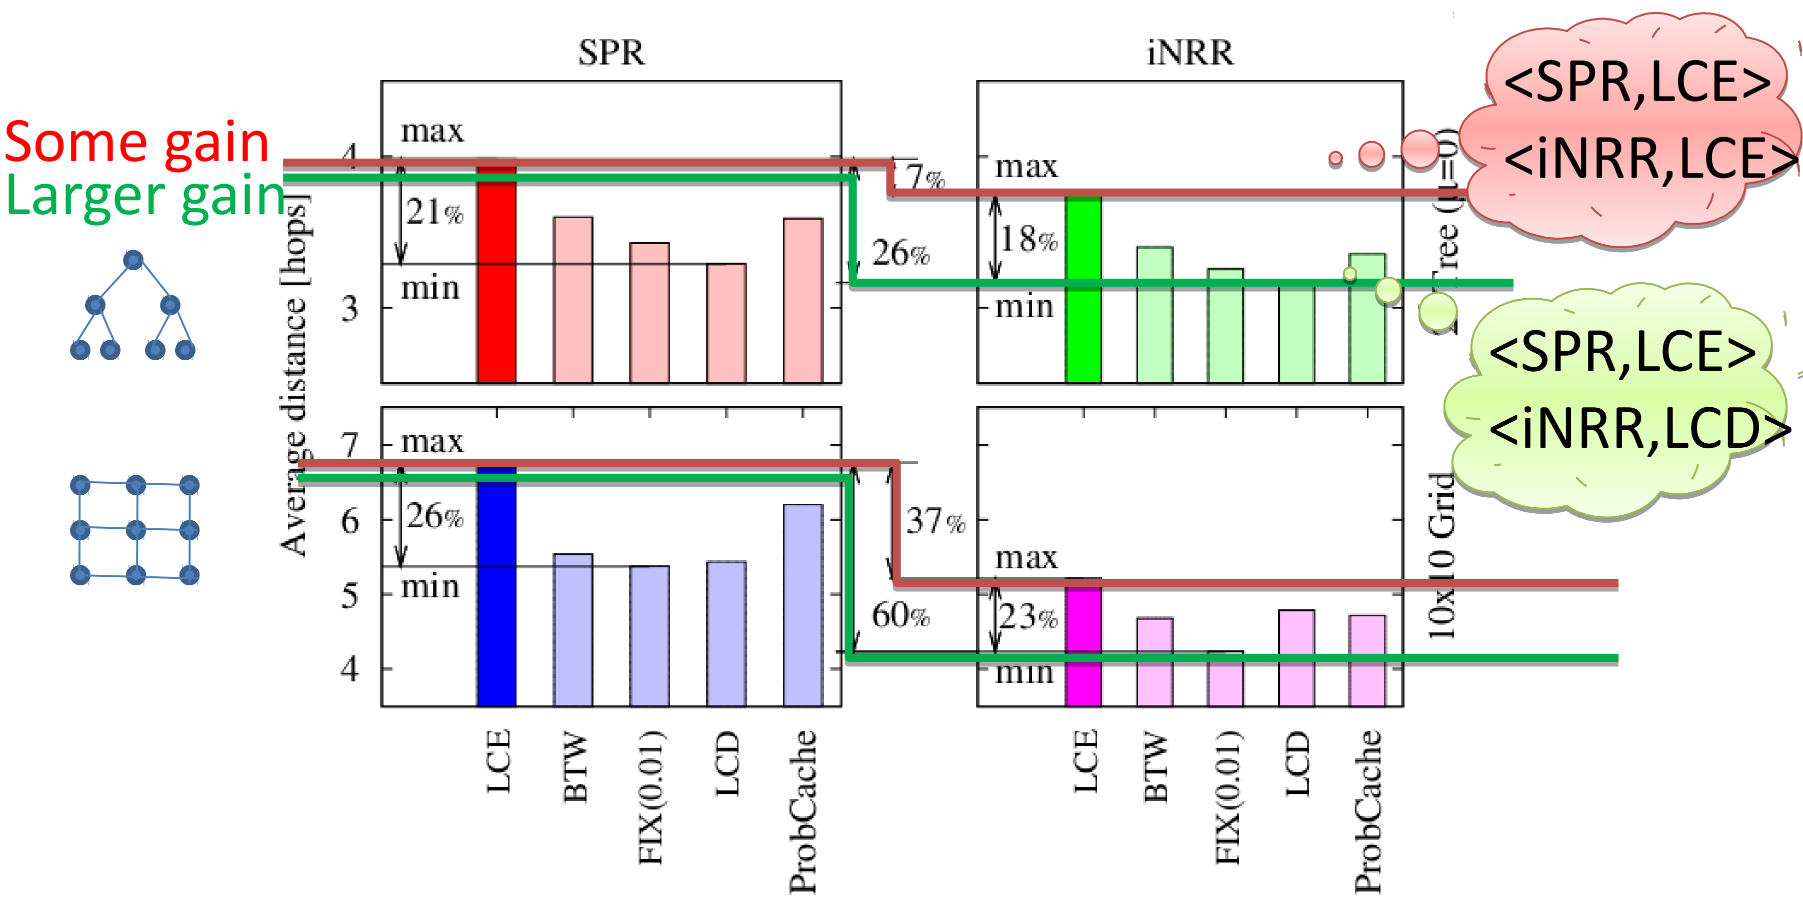

Fig. 2. ⟨F,D⟩ performance at a glance: average content distance as a function of meta-caching policies, for SPR (left) and iNRR (right) forwarding, on tree (top) and grid (bottom) topologies.

More realistic catalog/cache scales