Collaborative Caching in Content Networking

Liang Wang

Supervisor: Prof. Jussi Kangasharju

Email: liang.wang@cs.helsinki.fi

Department of Computer Science

University of Helsinki, Finland

About Me

Specialized in collaborative algorithm design for content networking, modeling and analysis of communication protocols in complex systems.

I am also interested in sensor networks, energy-efficient and green networks, social network analysis, and opportunistic networks.

- Education
  - B.Eng., “Computer Science & Mathematics”, in Tongji University, China.
  - M.Sc., major in “Computer Science”, minor in “Mathematics”, in University of Helsinki, Finland.
  - Ph.D. Candidate, “Computer Networks”, in University of Helsinki, Finland.

Content

Today's talk will cover the following topics

- Background of content networking
- Evaluation methodology for cache networks
- Model of collaboration
- Optimization and cost analysis
- Content discovery and delivery
- Kvasir project

Why Do We Need Content Networking

- Content distribution is the primary task for today's Internet.
- Traditional paradigm of communication network is Point-to-Point.
- Point-to-Point paradigm has many drawbacks when dealing with large-scale content distribution - efficiency, security and privacy.

Content consumer only cares what it is instead of where it is from.

Information-Centric Network Architecture

Many proposals exist in the literature, but we focus on ICN.

ICN is a clean-slate redesign of the current Internet infrastructures,

- Content is accessed by name.
- Caching is universal in the network.

ICN tries to solve the problems confronting the current Internet, e.g., content distribution efficiency, security, network congestion and etc.

Meanwhile, ICN poses new challenges on cache management, content addressing, routing and etc.

Cache Network Model

Given a group of networked caches, how to utilize them smartly and efficiently in order to push the system to its optimal state?

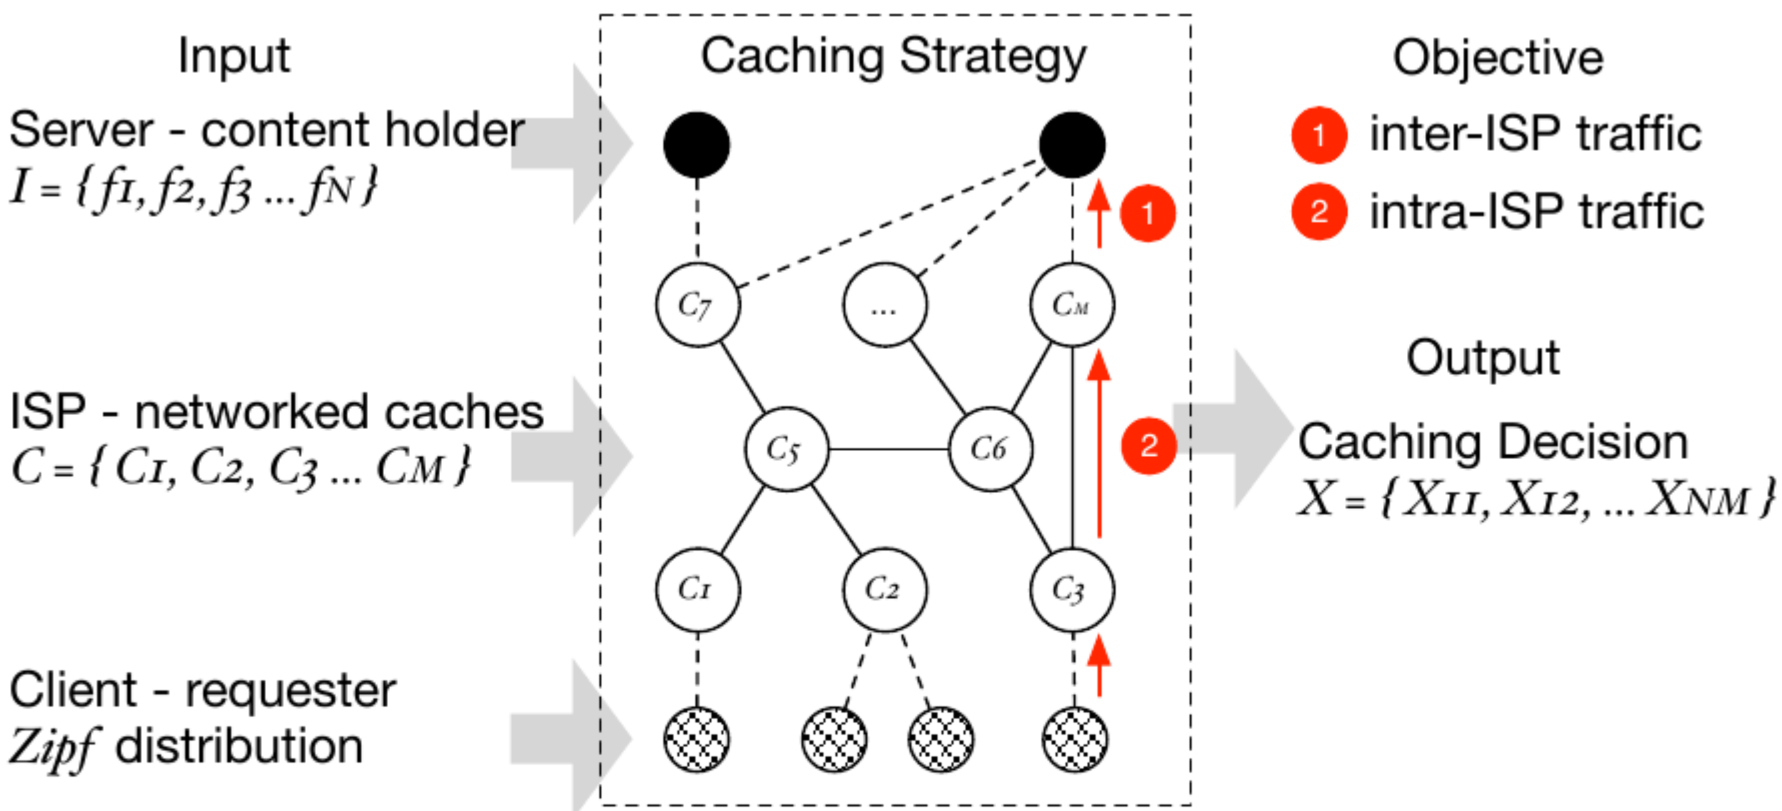

Essentially, we need manage a group of networked caches ….

Cache Network Model

Given a group of networked caches, how to utilize them smartly & efficiently in order to push the system to its optimal state?

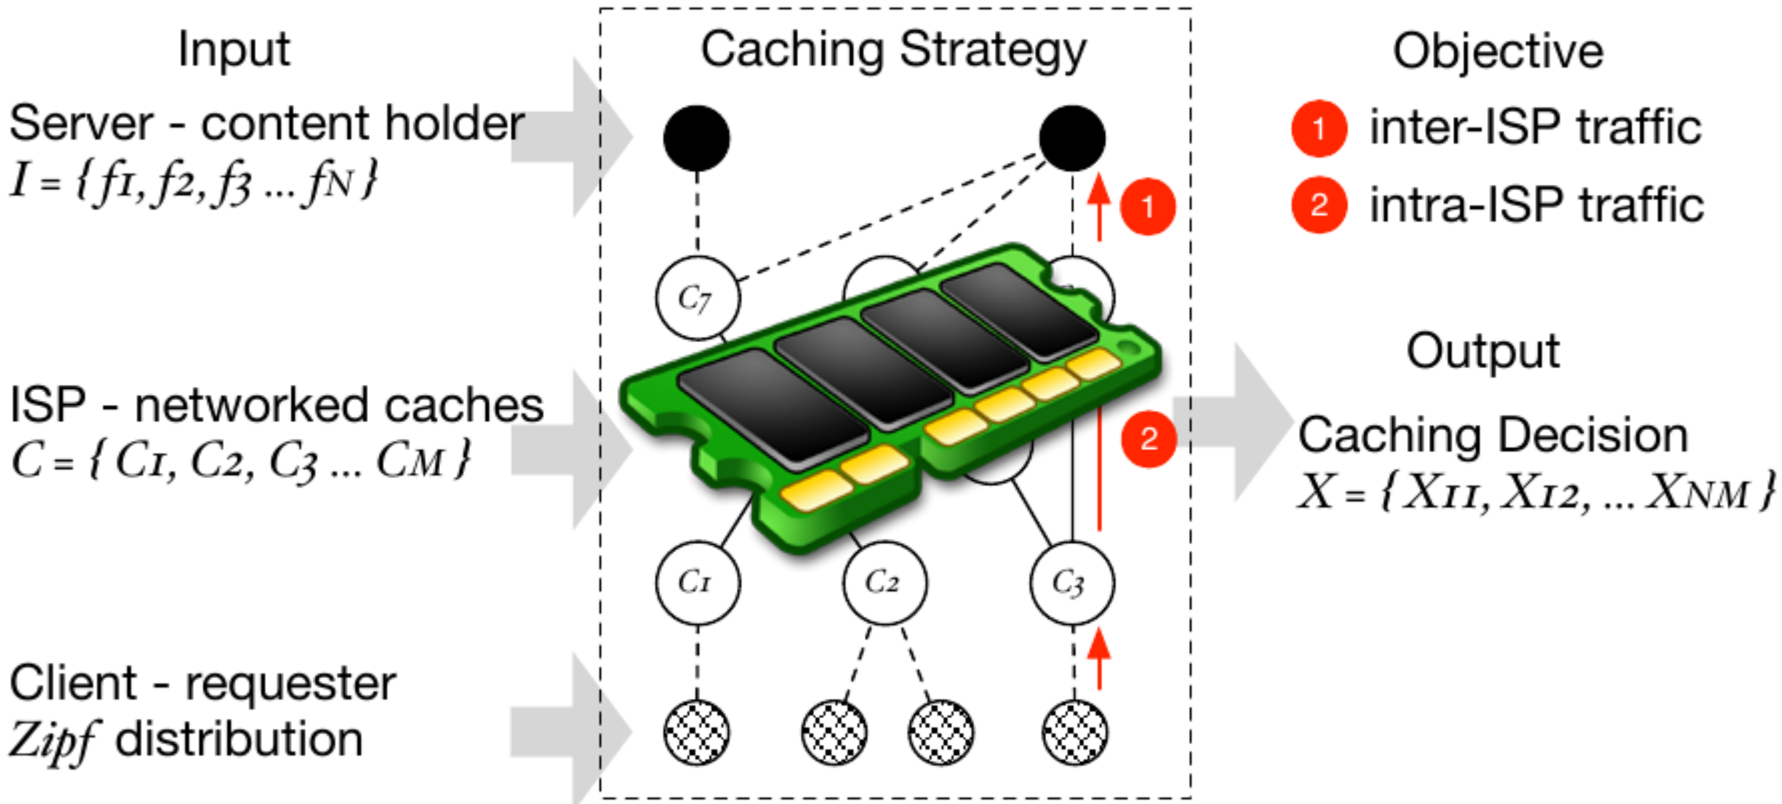

We want to use them as a single big cache ….

Simple Question, But Hard Challenge

We mentioned “the optimal state of the system”,

But wait ....

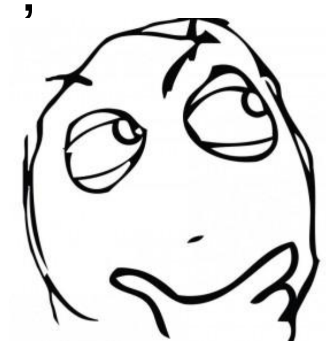

- How do you define the optimum?
- What metrics you use to quantify the performance?
- How are you going to evaluate an ICN design?

We need enough metrics to build up a holistic view of ICN systems.

Single Cache vs. Cache Network

What is the fundamental difference between a single cache and a cache network? Network structures.

- Content caching ≠ Content addressing
- Effective capacity ≠ Aggregated cache size
- Local optimum ≠ Global optimum

The system should not be treated as a simple "entity", we need zoom in to have a better look at its internal network structure.

Metrics

Content networking concerns traffic localization

→ i.e. how much saving on inter-ISP and intra-ISP traffic?

→ i.e. how many hits and where they occur in the network?

Byte hit rate (BHR) - saving on inter-ISP traffic.

Footprint Reduction (FPR) - saving on intra-ISP traffic. latency and etc.

Coupling factor (CPF) - what is this then?

Metrics - Byte Hit Rate

- Byte hit rate (BHR)
  - Stone age metric, measures the performance of a single cache.
  - Old but still useful, but only provides very limited information.
  - In cache networks, measures the saving on inter-ISP traffic.
  - We call it 1-dimension metric, since it only tells you the HOW MANY hits happens in the WHOLE network.

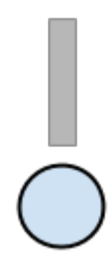

The whole network is treated as a single entity.

Metrics - Footprint Reduction

- Footprint Reduction (FPR)
  - FP is defined as the product of traffic volume and traffic length.
  - In cache networks, measures the saving on intra-ISP traffic.
  - We call it 2-dimension metric, since it tells you both HOW MANY hits and WHERE they happen along a path (source <-> destination) on average in the network.

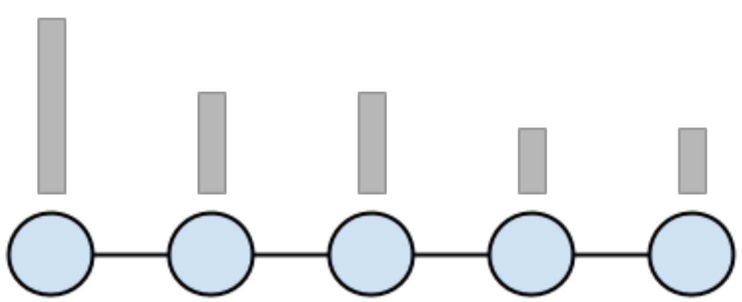

Performance is measured on per path basis.

Metrics - Coupling Factor

- Coupling Factor (CPF)
  - Defined as Pearson correlation between the popularity of the cached content and node's betweenness centrality.
  - Provides more information than both BHR and FPR.
  - We call it 3-dimension metric, since it tells you HOW MANY hits and WHERE they happen in the NETWORK (e.g., edge or core).

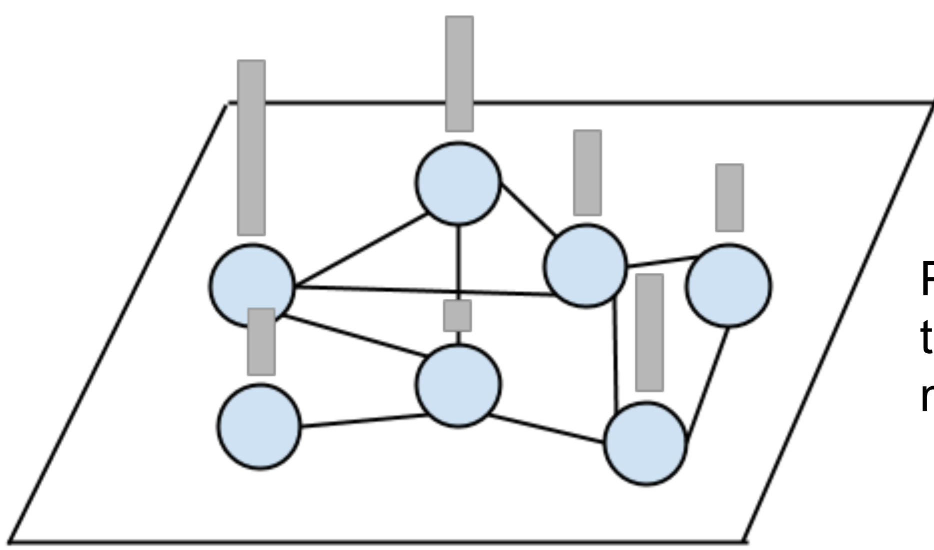

Performance is measured and correlated with their position in the network. The position is measured with betweenness centrality.

Metrics - Summary

Three metrics, are these all we need?

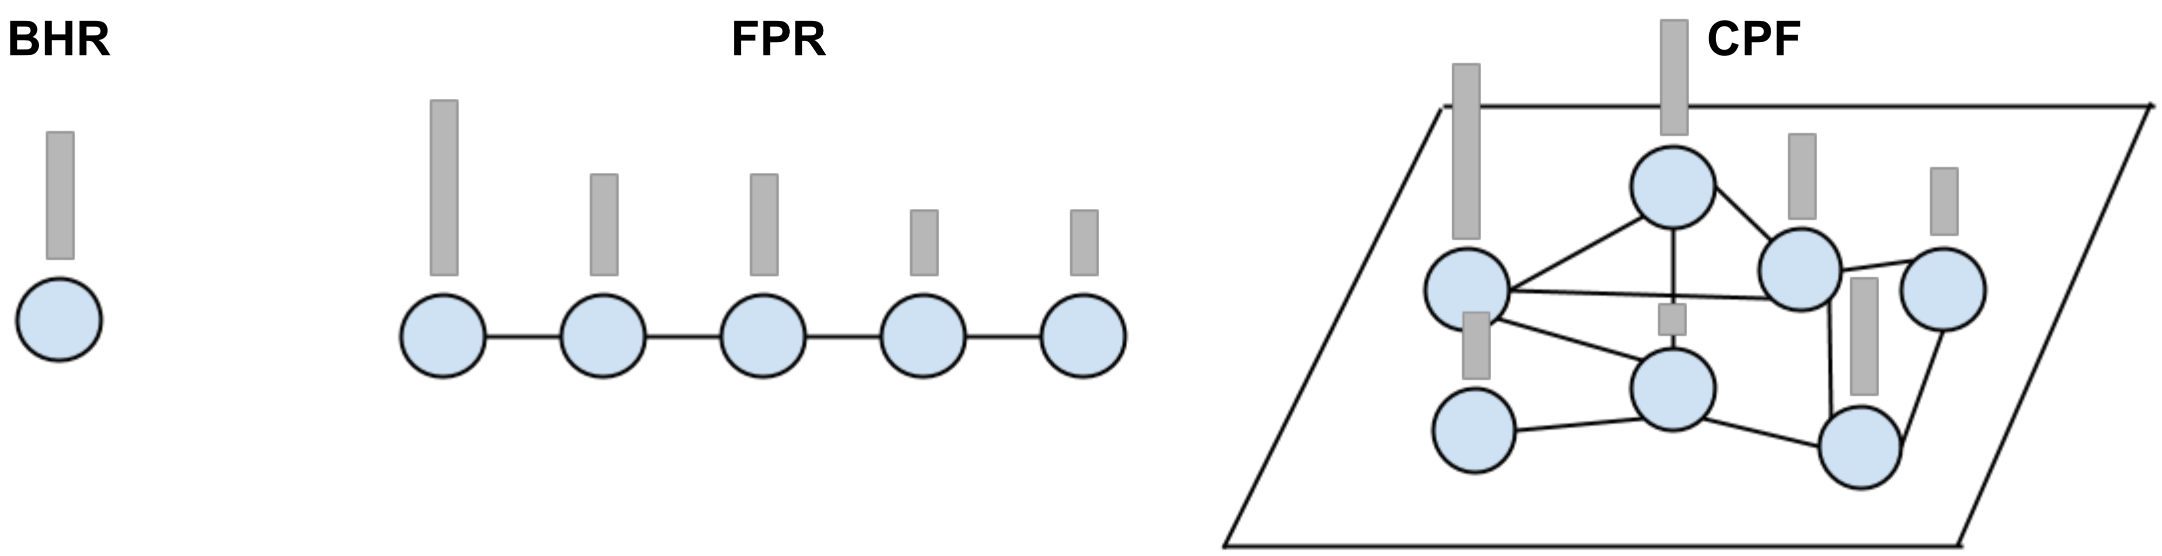

This set of metrics is far from complete, but it is sufficient for today's discussion on cache networks. E.g., the 4th dimension time can be included to model aging and evolution.

How do they work in an actual evaluation?

The Role of Collaboration

In conventional single cache context, admission control and replacement policy answers WHAT question.

- What content to cache?

In cache network context, collaboration answers WHERE question.

- Where to cache the popular content?
- Where to fetch the popular content?

W. Wong, L. Wang, and J. Kangasharju, ”Neighborhood Search and Admission Control in Cooperative Caching Networks,” in the Proceedings of IEEE Globecom. IEEE, December 3-7 2012.

Two Diametrically Opposing Viewpoints

Negative view

Completely a waste of time, the cache should only be deployed at the network edge.

Low utilization in upstream caches due to the strong filtering effect.

…..

Technique:

Analytical model, simulation, optimization and etc.

Positive view

Even naive collaboration can boost the caching performance.

Storage price keeps dropping, cache is pervasive nowadays.

….

Technique:

Analytical model, simulation, optimization and etc.

Why different conclusions?

Two Diametrically Opposing Viewpoints

Negative view

Completely a waste of time, the cache should only be deployed at the network edge.

Low utilization in upstream caches due to the strong filtering effect.

…..

Technique:

Analytical model, simulation, optimization and etc.

Positive view

Even naive collaboration can boost the caching performance.

Storage price keeps dropping, cache is pervasive nowadays.

….

Technique:

Analytical model, simulation, optimization and etc.

Different assumptions! Different metrics!

Two Diametrically Opposing Viewpoints

Negative view

- Regular tree structures.
- Requests come from leaves.
- Routers use simple LRU.
- Practically no collaboration.
- Usually just use BHR metric.

Positive view

- General topologies.
- Request can be pervasive.
- Various admission controls.
- Different level of collaborations.
- Usually use BHR and FPR.

Advocates of collaborative caching usually ask:

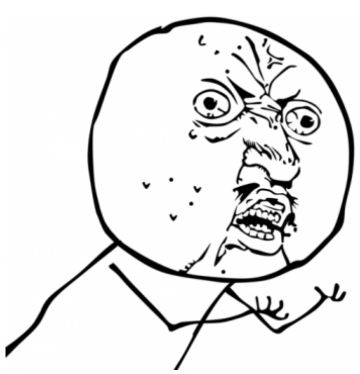

Why You No Use Those Networked Caches Like A Single Big Cache?!?!

Model of Collaboration

(K, r)-Collaboration Model

- r is the maximum search radius of a given node, it uniquely defines a neighborhood for collaboration. I.e., the range of collaboration.

- K is the maximum number of content replicas in the neighborhood defined by search radius r. I.e., the tolerance on duplicates.

L. Wang, S. Bayhan, and J. Kangasharju, "Effects of Cooperation Policy and Network Topology on Performance of In-network Caching," IEEE Communication Letters. IEEE, Vol.18, No.4, April 2014.

Rationale Behind K and r

(K, r)-Collaboration Model is simple yet expressive.

The rationale behind K and r

- We can always characterize collaboration by its search strength (r) and capability of reducing duplicates (K).
- In other words, the ability of discovering content (r) and the ability of utilizing cache efficiently (K).
- We can see they represent the tradeoff between BHR and FPR.
- We will also see how the collaboration model impacts the content distribution in the network.

Interplay of Content, Topology and Collaboration.

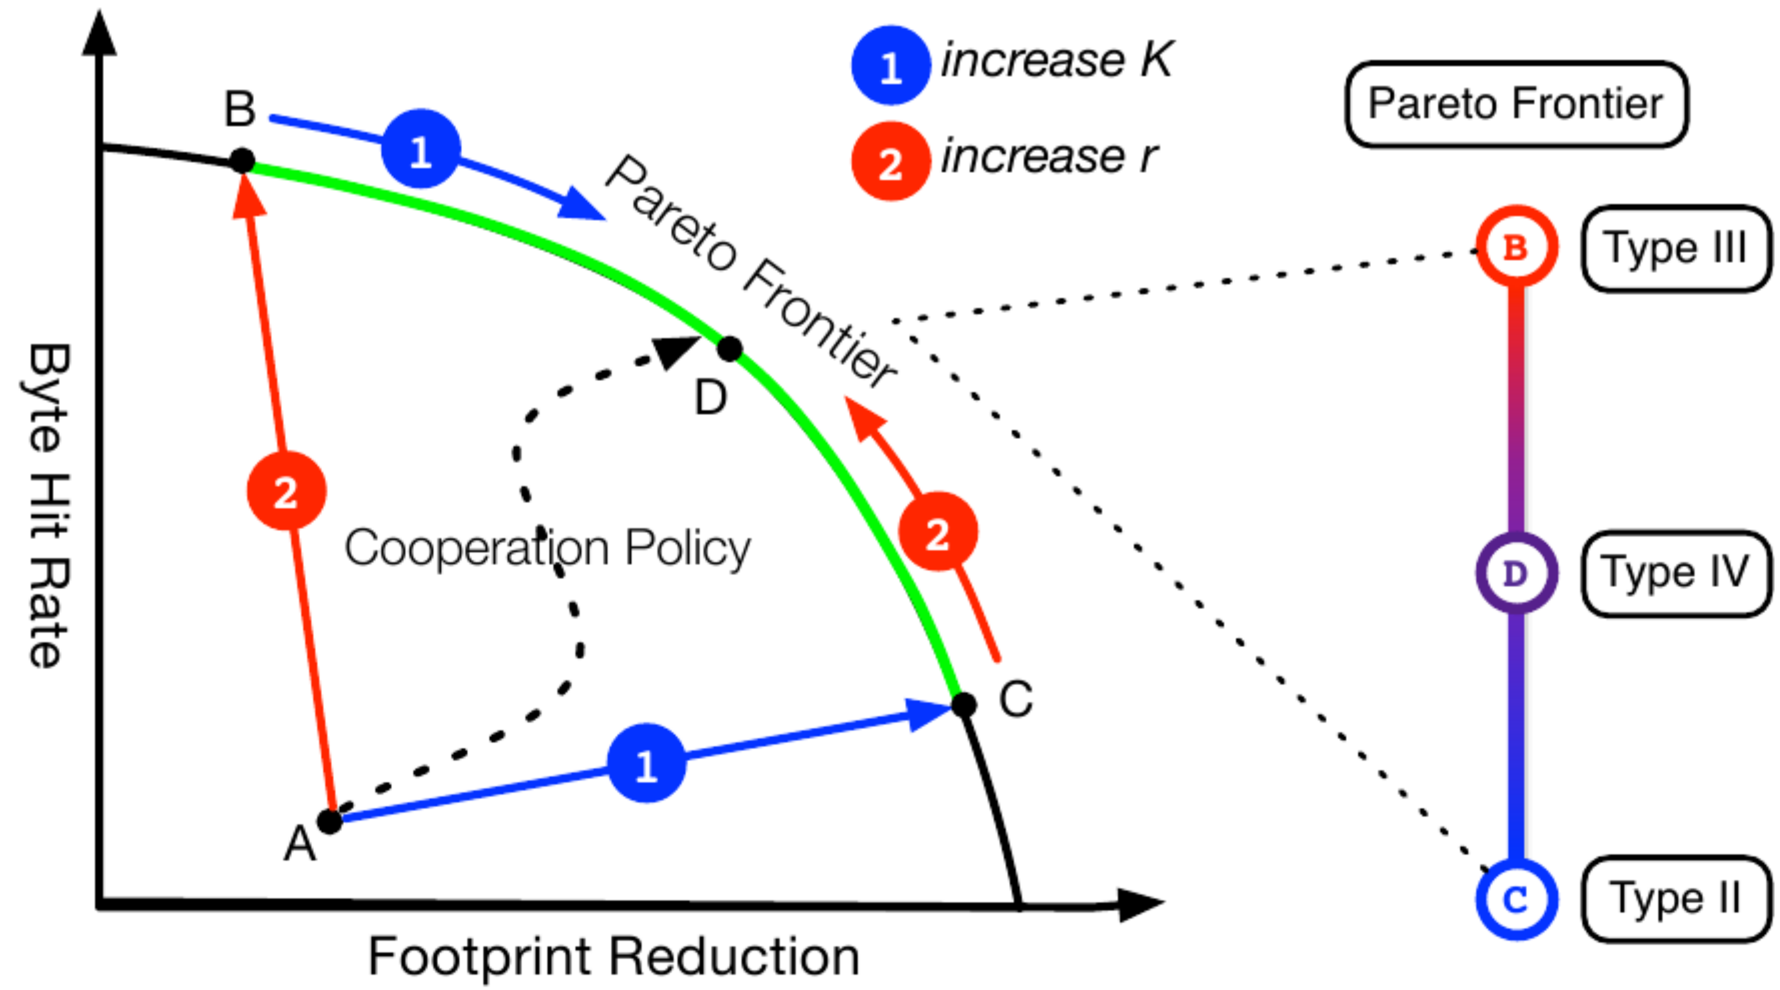

An important figure to describe the cache system behavior by using the metrics and collaboration model we presented in the previous slides.

Interplay of Content, Topology and Collaboration.

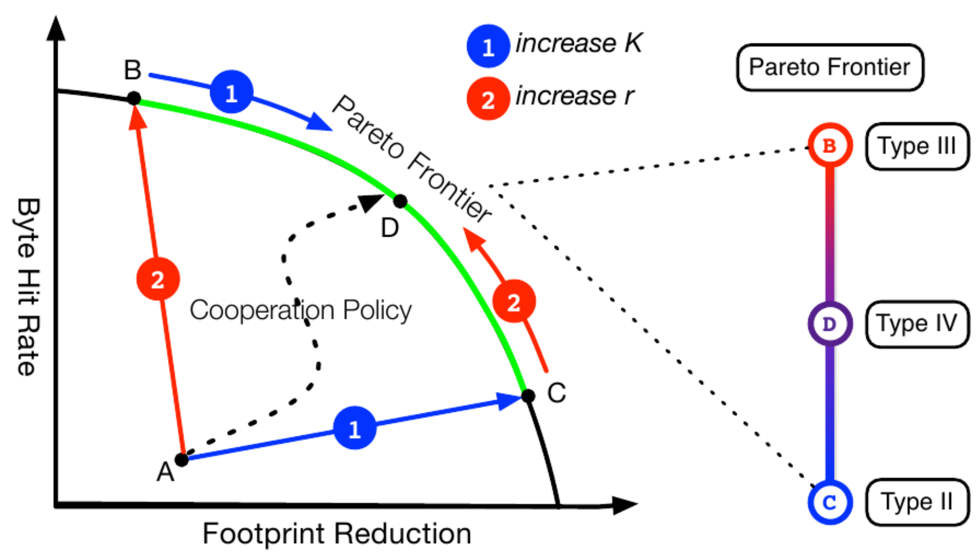

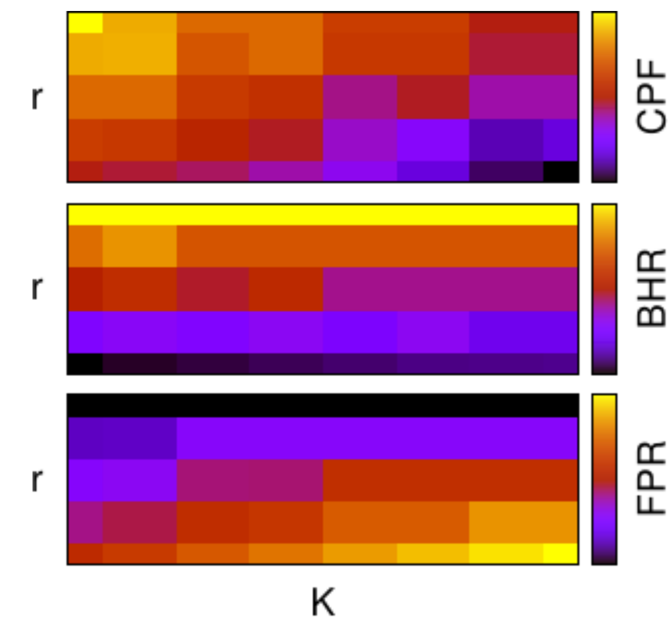

Now, let’s take a tour A → B → D → C → A

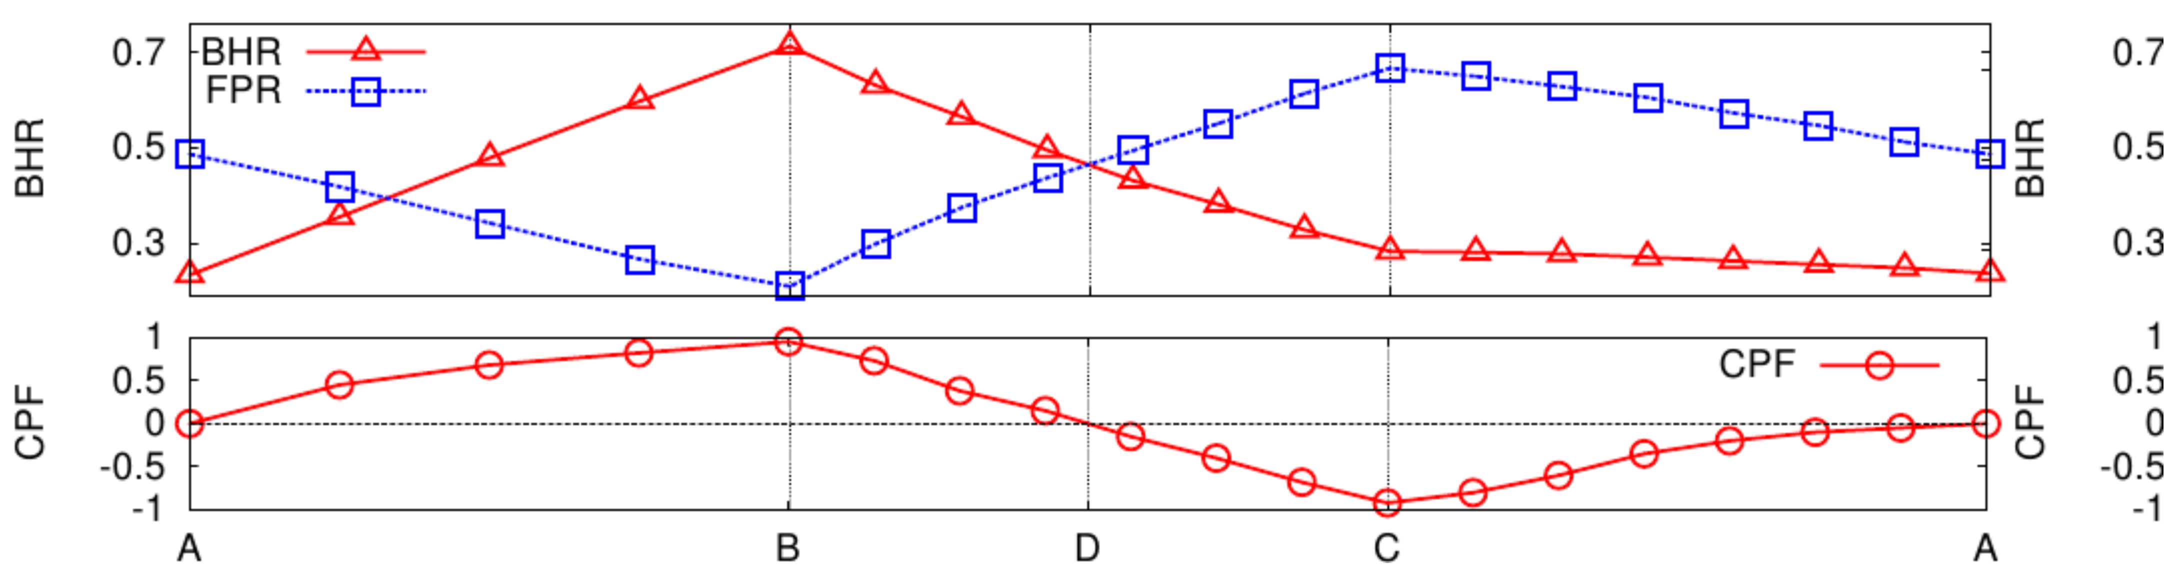

Interplay of Content, Topology and Collaboration.

Let's zoom in the interesting Pareto frontier (B→D→C) where CPF varies in (-1, +1). What does it mean in practice? Recall our definition.

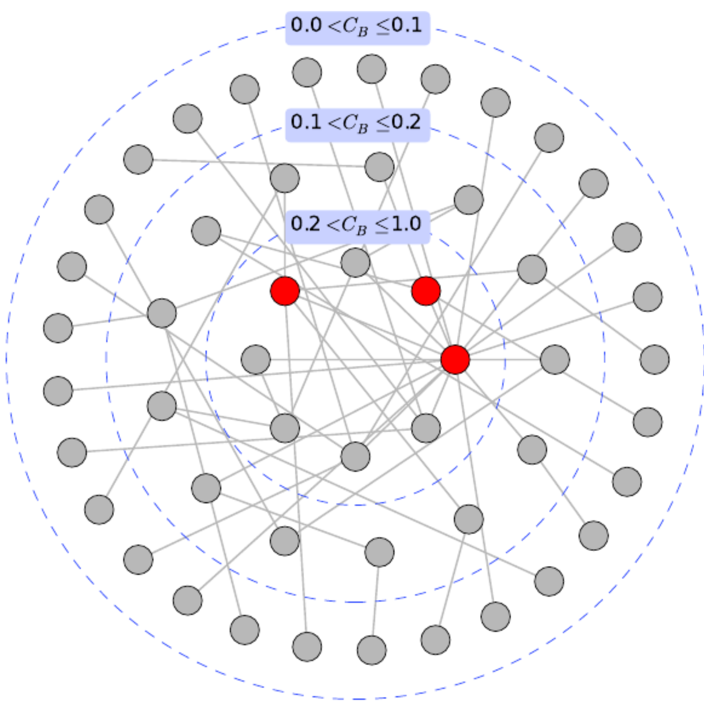

B

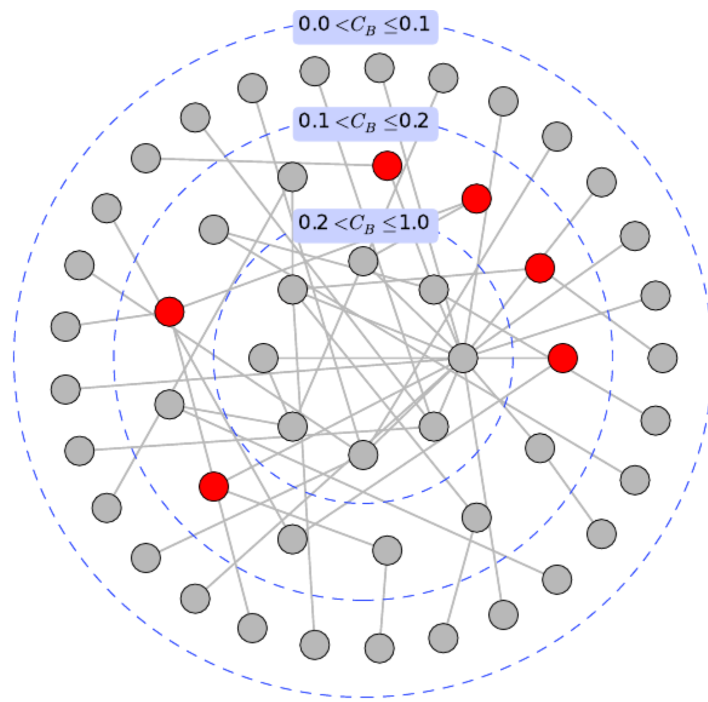

D

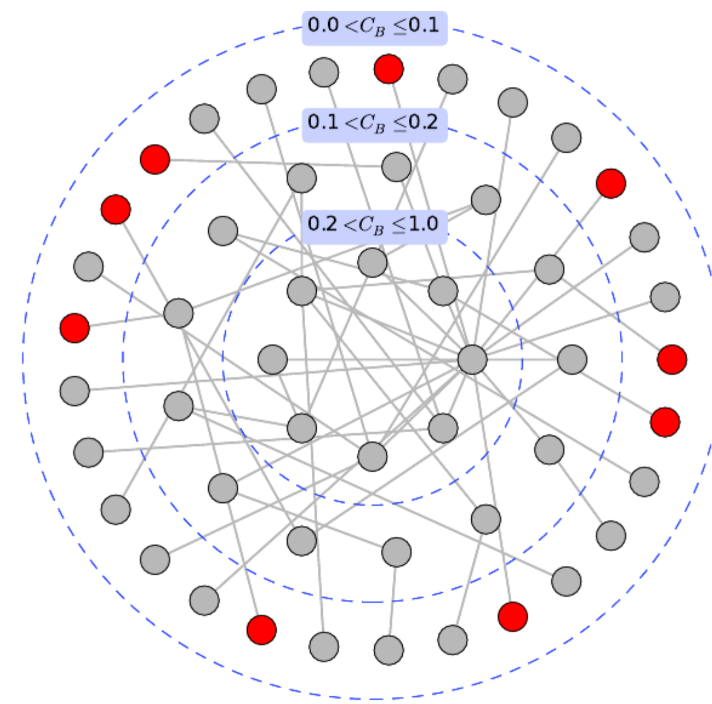

C

Collaboration glues/couples content with topology.

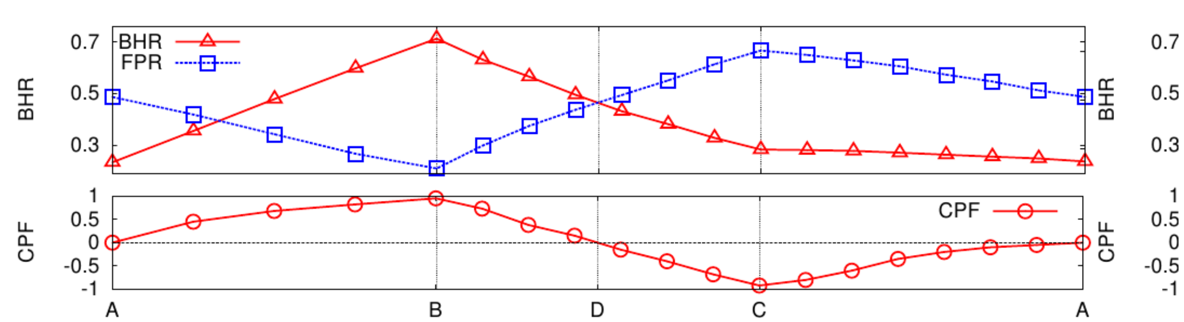

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Beyond BHR and FPR

- We have potentially infinite Pareto optimal solutions on the frontier.
- How are we going to find the point “D” on the Pareto frontier?
- Different tradeoff between BHR and FPR.

If all are optimal, then ...

- What collaborative strategy is considered as a good strategy?
- What have we missed in our model?

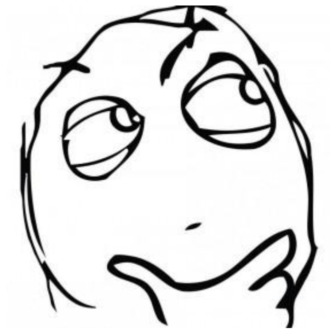

The Basis of Collaboration

What have we missed in our model?

- Incentive for collaboration. We assumed nodes are altruistic.
- What if selfishness is an inherent and intrinsic characteristic?
- It is hard to justify that a node would like to sacrifice for others.
- What is the fundamental basis of being collaborative?
- Fairness is important!

L. Wang and J. Kangasharju, "A Fair Collaborative Game on Cache Networks," in submission.

Fair Collaborative Game on Cache Networks

What is the system model of a fair collaborative in-network caching game?

- Each node is associated with a utility function.
- Node gets rewards by satisfying the requests from the clients.
- The content can be cached locally or retrieved from neighbors.
- The benefit decreases if the retrieval distance increases.
- Nodes collaborate to determine what to cache and where to fetch.

Definition 3. *An in-network caching game is a tuple* (Ω, u^0)*, where* $\Omega \subset \mathbb{R}^{|V|}$ *contains all the utility values obtainable via collaboration,* $u^0 \subset \mathbb{R}^{|V|}$ *contains all the disagreement values leading to a negotiation breakdown.*

Nash Bargaining Framework

How to find a efficient and fair solution for the game?

- Formulate the problem in Nash Bargaining framework.
- Axiomatic game theory, agnostic about negotiation mechanisms.

Definition 4. *A fair collaborative game is a game* (Ω, u^0) *with Nash bargaining solution, namely a function* $f : \Omega^e \to \Psi$ *such that* $f(\Omega, u^0) = (\mathbf{x}, \mathbf{y})$ *uniquely maximizes* $\prod_{v_i \in V}(U_i - u_i^0)$.

Solve the Optimization Problem

Relatively straightforward convex optimization techniques.

Boils down to solving a huge and tedious equation system.

Good news is that we can decompose the equation system using

Lagrange Dual Decomposition, then break it down to neighborhood size.

- Not only provides us a decentralized solution.
- But also exposes the structure of collaboration.
- Shadow price for exchanging content is the complicating variable.

Cost Analysis of Collaboration

Instead of trivial topology with regular structure like a line or tree, we are more interested in the collaboration cost on general topologies.

- The cost is measured in terms of number of exchanged messages.
- The cost grows exponentially when the search radius increases.

Theorem 3. *In a random network $G = (V, p)$ where nodes have average search radius r, the induced system overhead $\Delta_r^{r+1}\Phi$ by increasing the average search radius by 1 equals*

$$\Delta_r^{r+1}\Phi = \theta \times |V| \times \left[\frac{z_2}{z_1}\right]^r \times z_1 \qquad (16)$$

Cost Analysis of Collaboration

Important implication

Theorem 3 conveys an important message on collaborative caching, and shows that the collaboration cost grows exponentially on most natural graphs like Internet and ISP networks. Therefore, the collaboration has to be restricted to a very small neighborhood to keep cost reasonable.

Collaboration is so expensive, are we doomed?

Collaboration Localization

Is collaboration doomed due to its cost? Fortunately, there is hope -- collaboration localization. See the results on Sprint network.

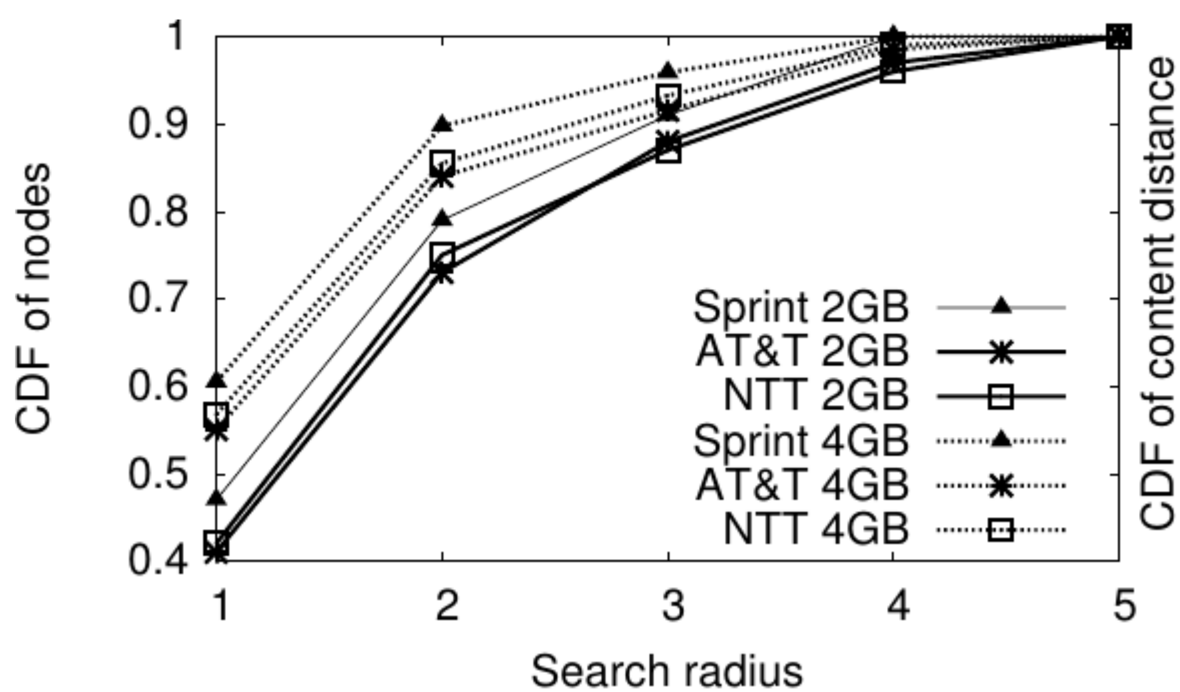

(a) CDF of nodes vs. search radius.

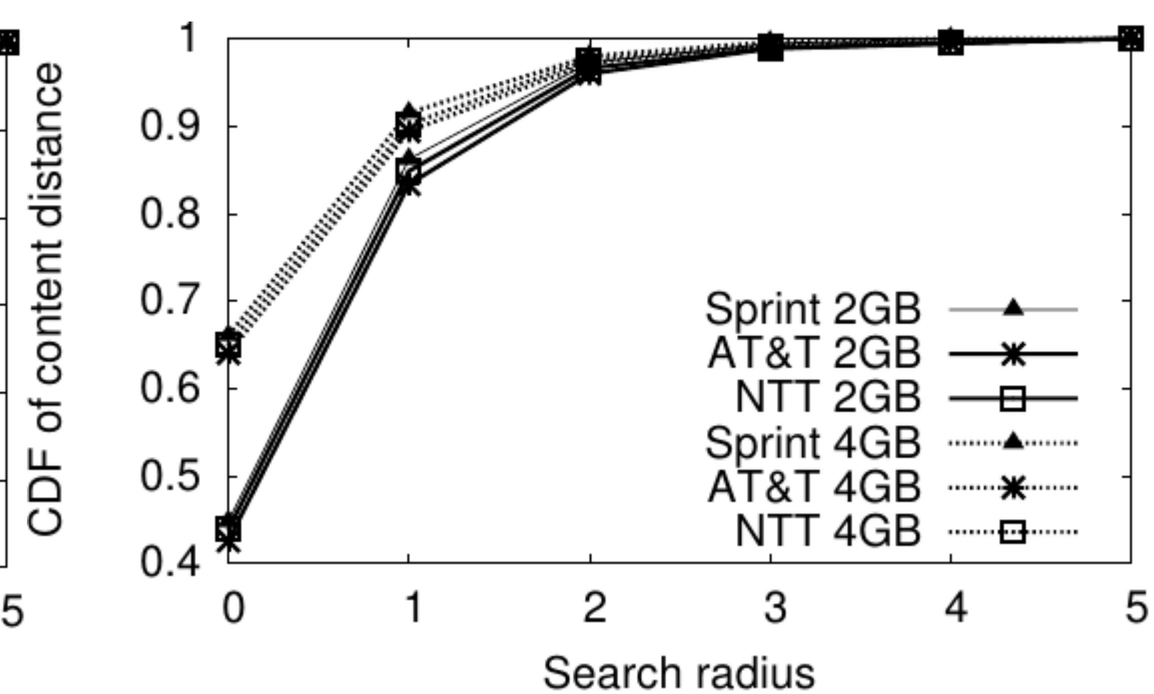

(b) CDF of content vs. search radius.

In practice, the optimal neighborhood is very small and most content is retrieved from the neighbors within 2 hops.

Collaboration Localization

Further investigation strongly indicates the collaboration is highly localized in a small neighborhood due to the highly skewed content popularity distribution.

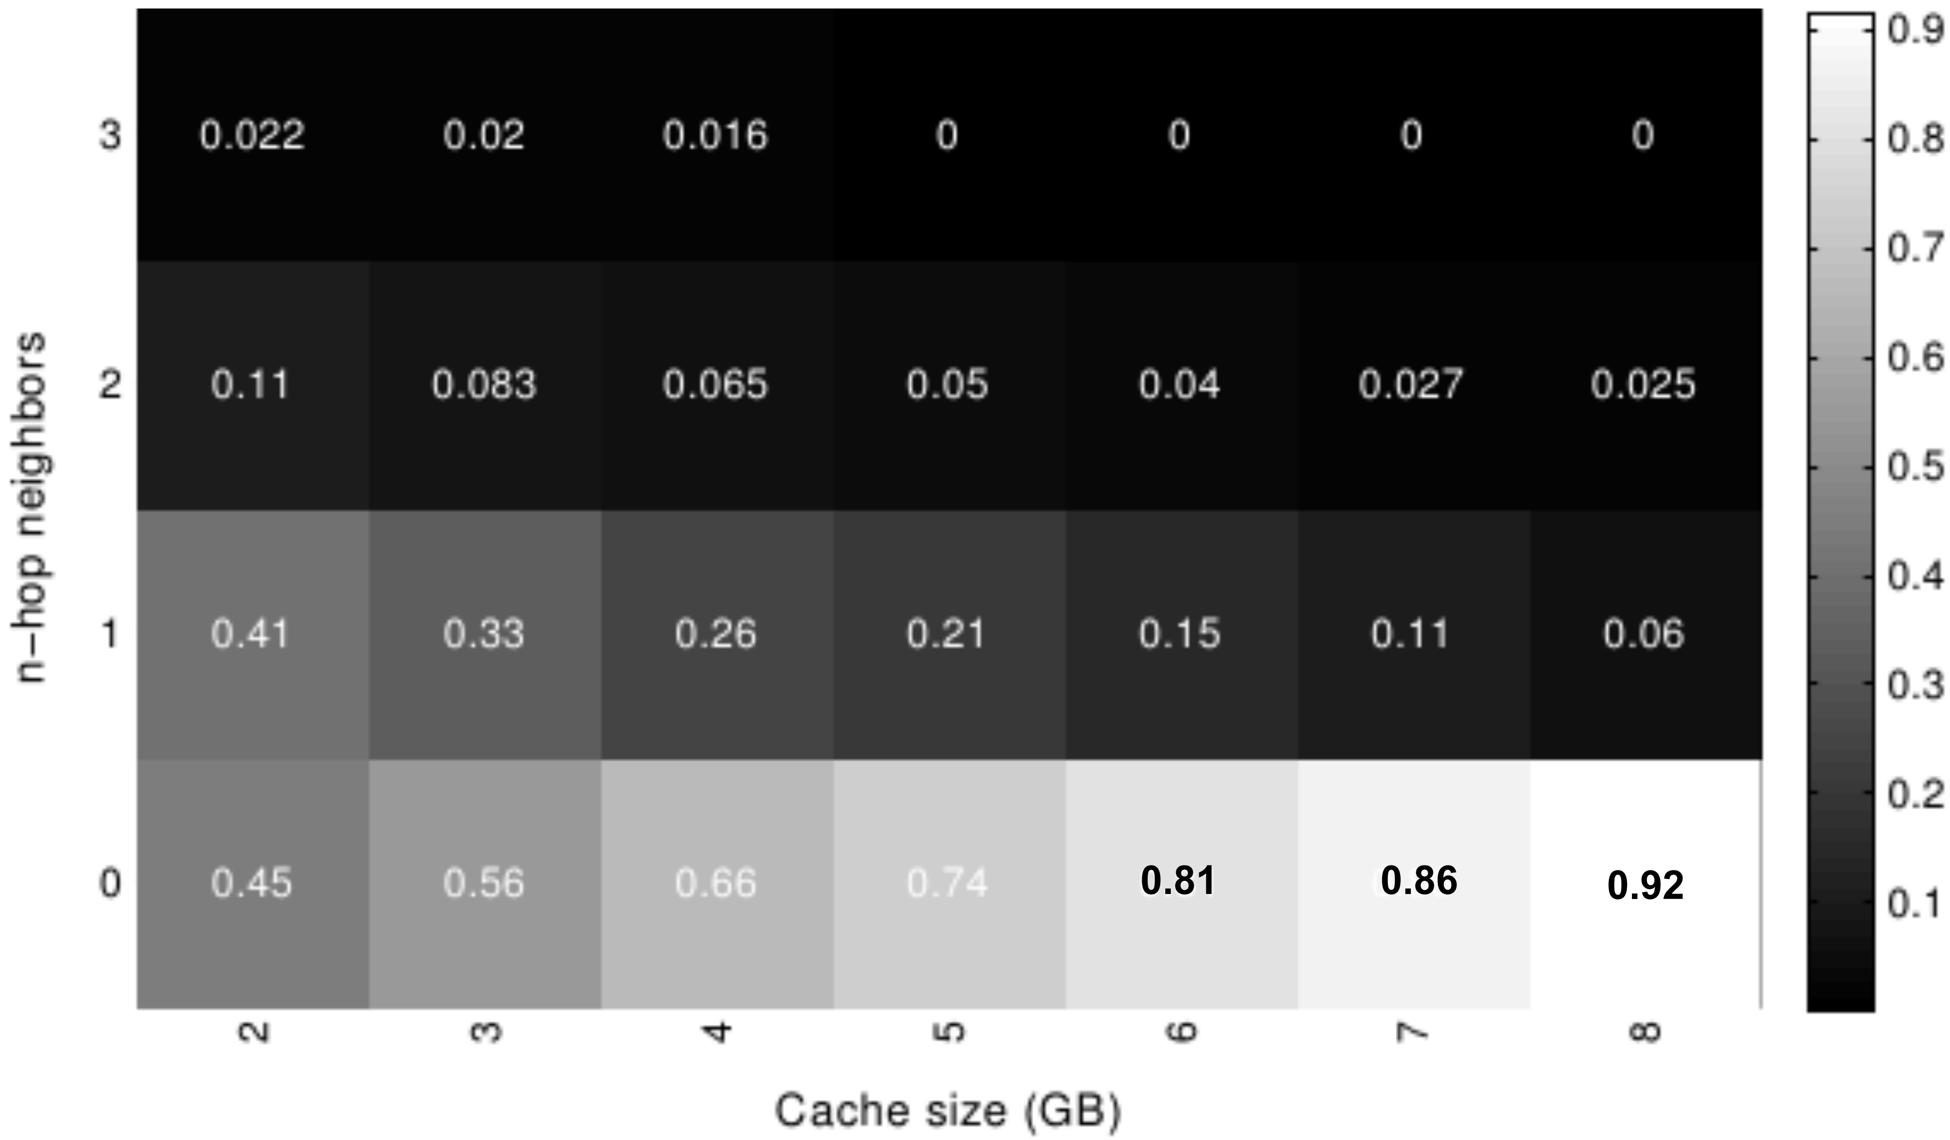

Well-Defined Fairness Metrics

Three well-defined fairness were investigated.

- Egalitarian Fairness - pursues the absolutely same amount of improvement on each node, and usually leads to a Pareto inefficient solution.

- Max-Min Fairness - pursues the fairness which maximizes the node with the worst utility.

- Proportional Fairness - pursues both proportional improvement on all nodes and maximizing the utility from collaboration.

Fairness Achieved in the Game

Proportional Fairness also indicates the optimal strategy is the one which can maximize the aggregated utilities from collaboration. Any deviation from the optimal strategy is detrimental to the collaboration.

Proven in the paper that

- Proportional fairness is guaranteed by the solution.
- Whenever Egalitarian fairness is achieved, it also indicates Max-Min fairness.

How About Performance?

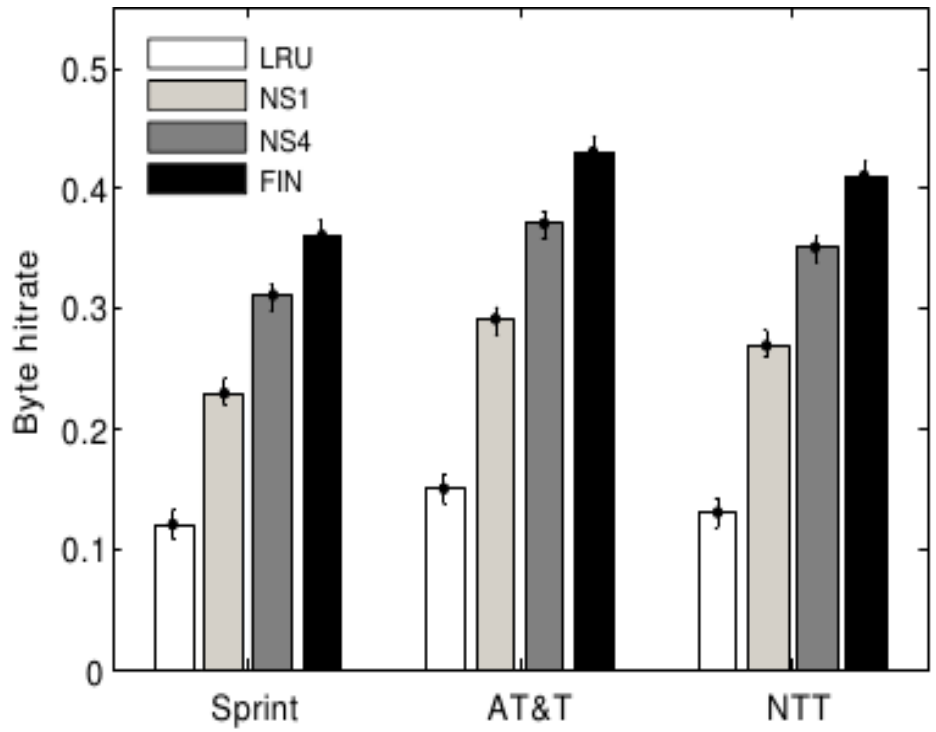

Byte hit rate comparison, 4GB cache.

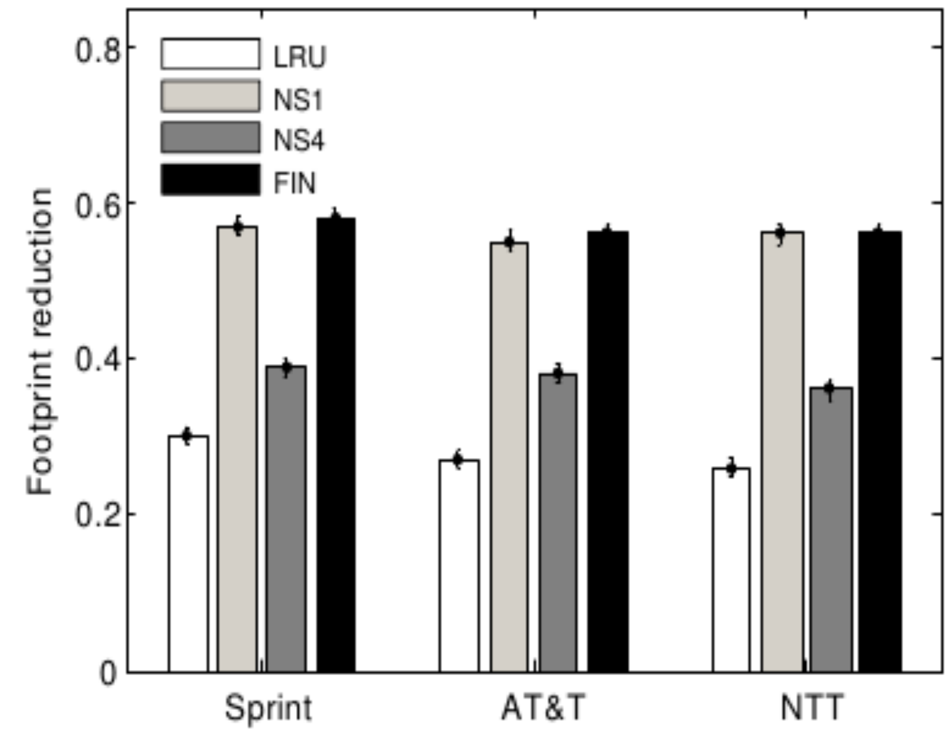

Footprint reduction comparison, 4GB cache.

LRU has the worst BHR and FPR, whereas our Fair In-Network caching algorithm (FIN) has the best. By increasing the search radius, NS4 achieves better BHR, but FIN consistently remains at least 16% better than NS4 over all the networks. Figure on the right shows NS4 has worse FPR (less than 40%) than NS1 and FIN, indicating the gain in BHR is achieved at the price of sacrificing FPR due to increased traffic.

Content Discovery and Delivery

So far our model only concerns the resources allocation, it does NOT specify how the content is discovered and delivered.

Namely, how the content should be addressed and how the query should be routed.

- Probabilistic solution.
- Deterministic solution.

L. Wang, O. Waltari, and J. Kangasharju, "MobiCCN: Mobility Support with Greedy Routing in Content-Centric Networks," in the Proceedings of IEEE Globecom. IEEE, December 9-13 2013.

Routing and Content Addressing

Some examples of different schemes

- The lightest solution is en-route discovery, simply forward the request to the next hop along the path. Lowest complexity, but no guarantees on finding the content.
- For small neighborhood, content discovery can be achieved by explicitly exchanging information on cached content. The communication cost can be reduced by using compression like Bloom Filter and etc.
- One extreme is search radius equals network diameter, then you have the whole network as your neighborhood. Highest BHR, but lowest FPR.
- Recall our collaboration model →

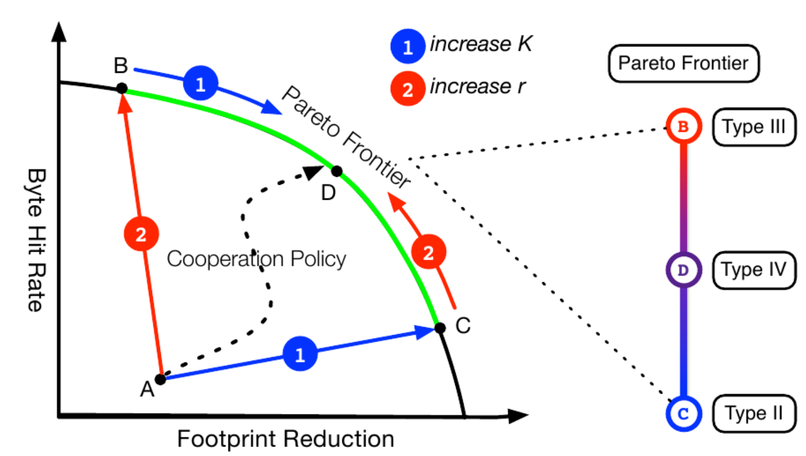

Greedy Routing Scheme

Greedy routing can be implemented as an underlay.

- Forward the packet to the closest neighbor (directly connected).
- Content addressing is usually done by hashing.

Pros

- Nodes have small routing tables.
- It solves scalability of name management issue.
- Routing protocol is very simple and low-complexity.
- Reduce the duplicate to the minimum level.
- Even solve the datasource mobility issue in content networks.

S. Roos, L. Wang, T. Strufe, J. Kangasharju, ”Enhancing Compact Routing in CCN with Prefix Embedding and Topology-Aware Hashing,” in the Proceedings of ACM MobiCom workshop on MobiArch. ACM, September 7-11 2014.

Greedy Routing Scheme

Cons

- Heavily relies on graph embedding procedure. (local minimum issue)
- Embedding procedure itself may not be scalable.
- May not be able to accommodate well to system dynamic.
- May have severe traffic and load balancing issue.

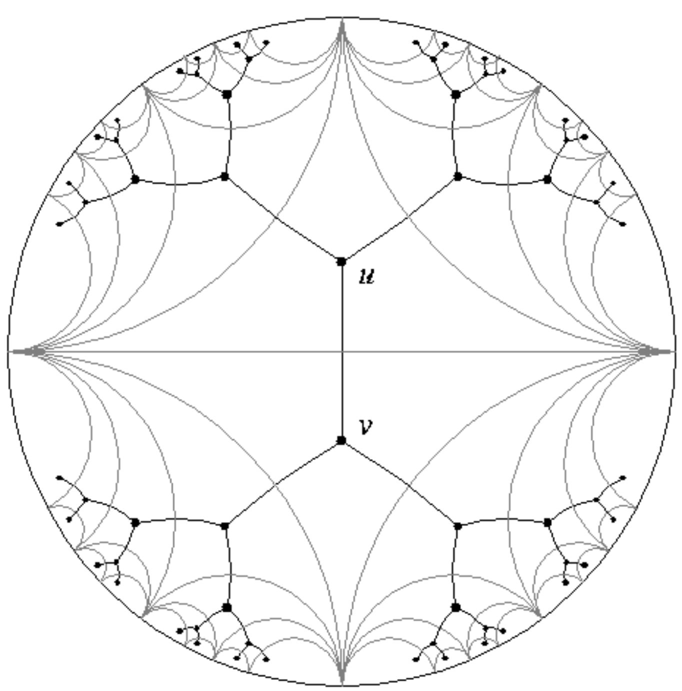

A example of hyperbolic embedding in Poincaré disc. It solves local minimum issue, but the root nodes suffer from imbalanced storage and traffic load.

Solution: Prefix-S embedding and topological-aware hashing.

Summary

- Collaboration on content networks can be modeled with (K,r) model.
- To gain a holistic view, measurement metrics must be carefully chosen.
- Given Pareto optimality, inter-ISP and intra-ISP are conflicting interests on non-trivial topologies.
- Fairness is the basis of being collaborative.
- The collaboration on general topologies is costly, it is only suitable when most of the gain can be obtained from a small neighborhood.

Acknowledgement

- Thank Bell Labs for the invitation, and hosting me in Dublin.
- Thank Alessandra for reviewing my application and initiating the visit.
- Thank Data Analytics Team (Alessandra, Deepak, Hamed, Marcel, Michael, Tiep, Yue) for spending time in interviewing me and their great patience and insightful questions.
- Thank Abhaya Asthana for reviewing and redirecting my application.
- Thank James for taking care of my VISA application.
- Thank Prof. Jussi Kangasharju and my group for their valuable comments on my work.

Q & A

Thank You!

Questions?

Prospect - Kvasir Project

The research focus should move from network layer to application layer.

Again, how do you understand ICN?

- Users only care what they need, not where they are from.
- User behavior has a clear pattern, not random noise.

Content objects are connected with their semantic meanings.

Maybe everything is already there, just need a little bit innovation to put them together! Kvasir Project.

Kvasir Project - Backend

The core is a highly optimized and high performance semantic engine.

- The research on the Backend brings up a lot of interesting challenges.
- How to efficiently reduce the dimensionality? SSVD, random projection ….
- How to perform fast search in high-dimensional space?
- How to effectively reduce the index size?
- How to build a better recommender based on one-class SVM?
- ……..

The backend is very fast and scalable comparing to the existing tools.

Kvasir Project - Frontend

Resides in the browser (as a plugin), connects the web pages with their semantic meaning. Alpha version is already fast and smart. :)

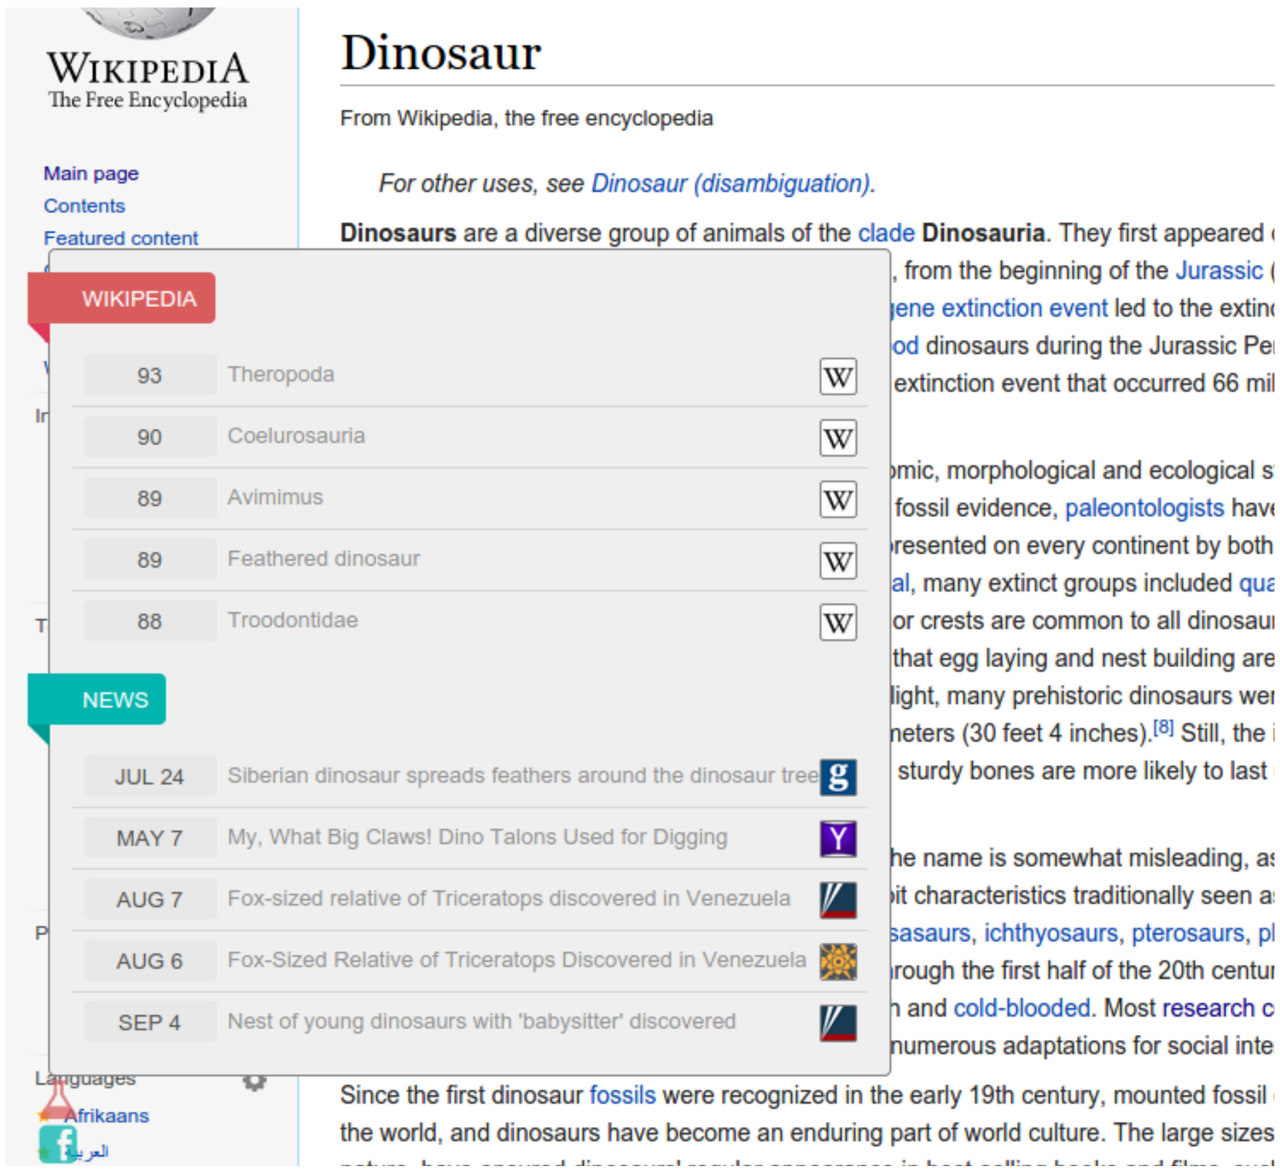

Provide the most relevant articles from various high-quality new sources and Wikipedia. The retrieval time is only 20 - 60 ms

Kvasir Project - More Examples

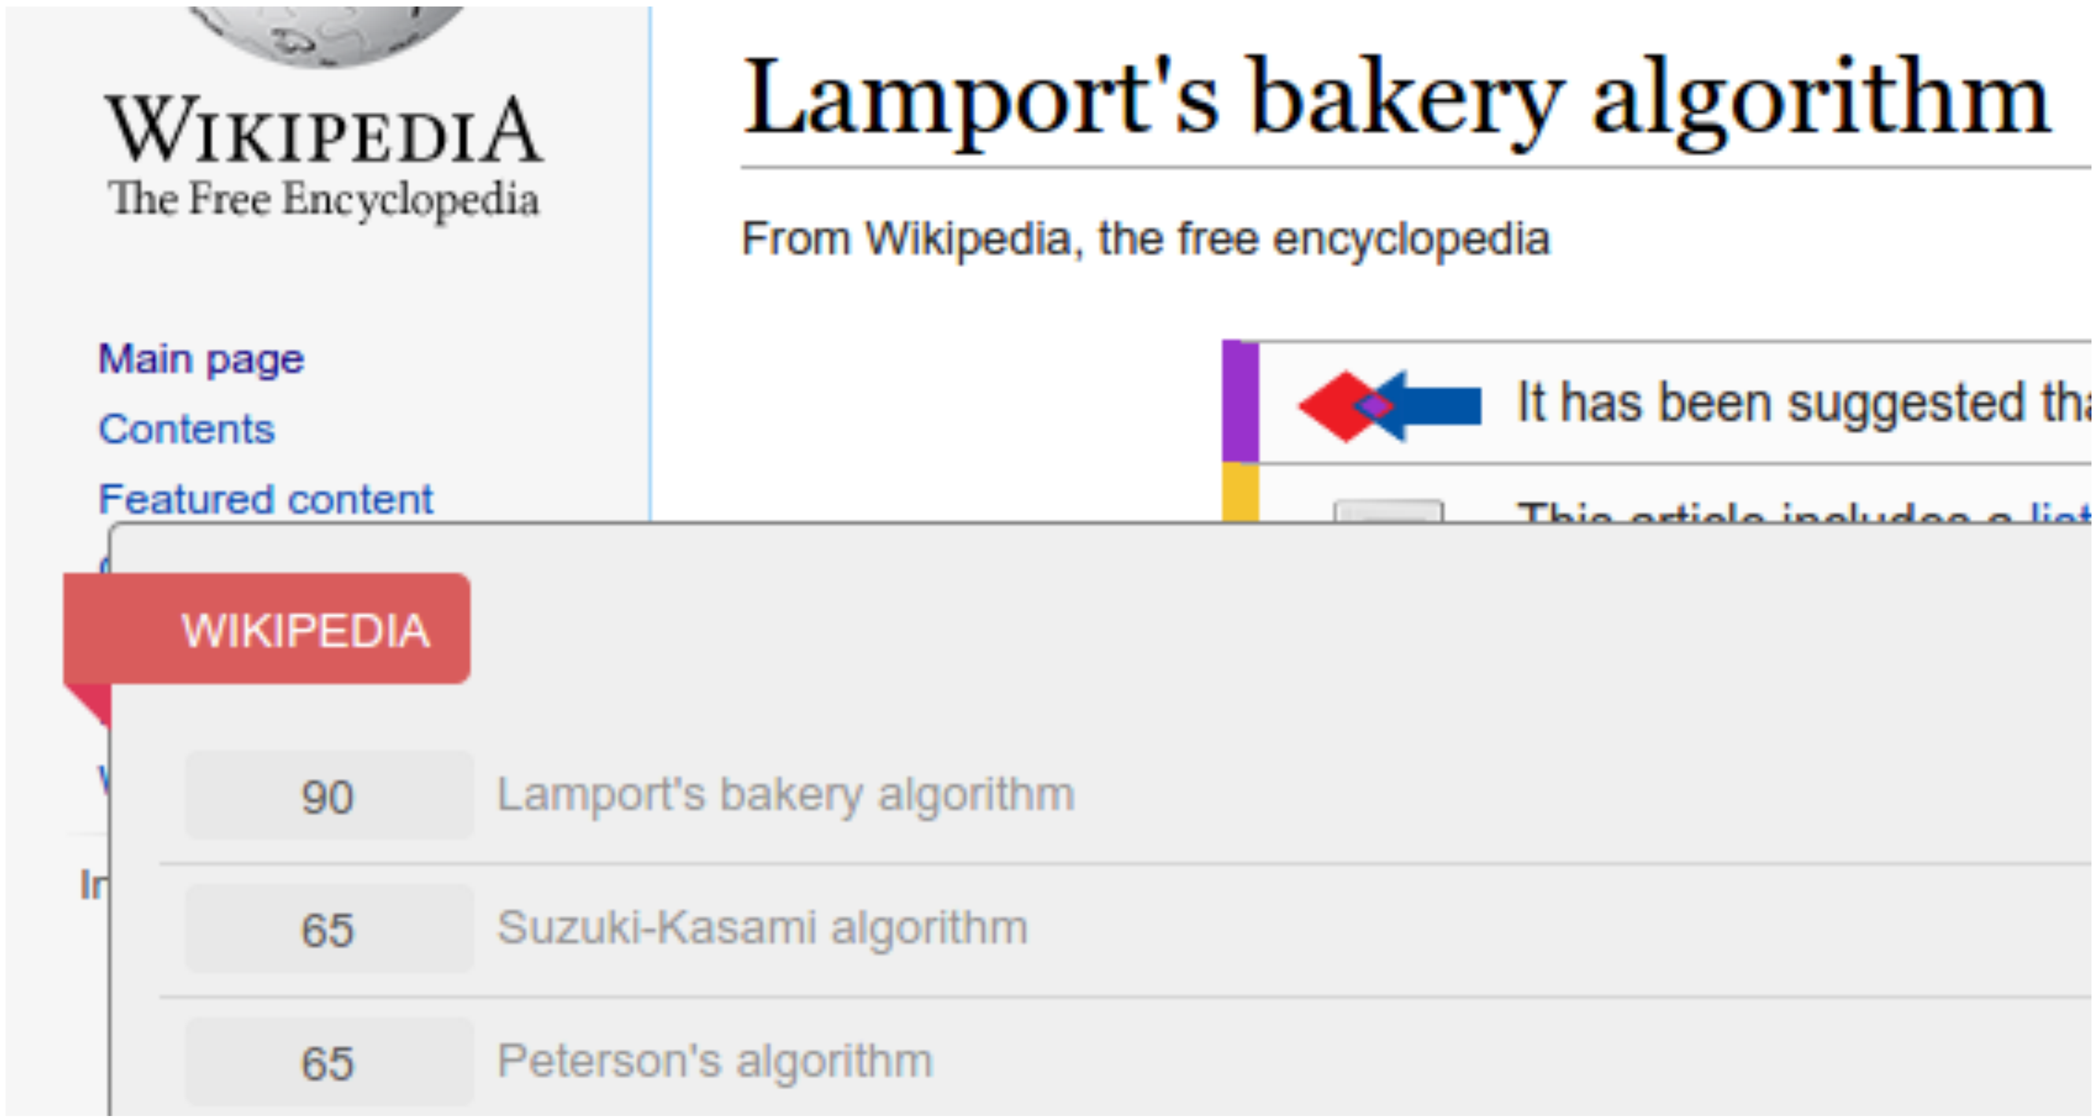

Note that Suzuki-Kasami algorithm is EITHER mentioned NOR appears on the page. But the Kvasir still successfully found it.

Kvasir Project - More Examples

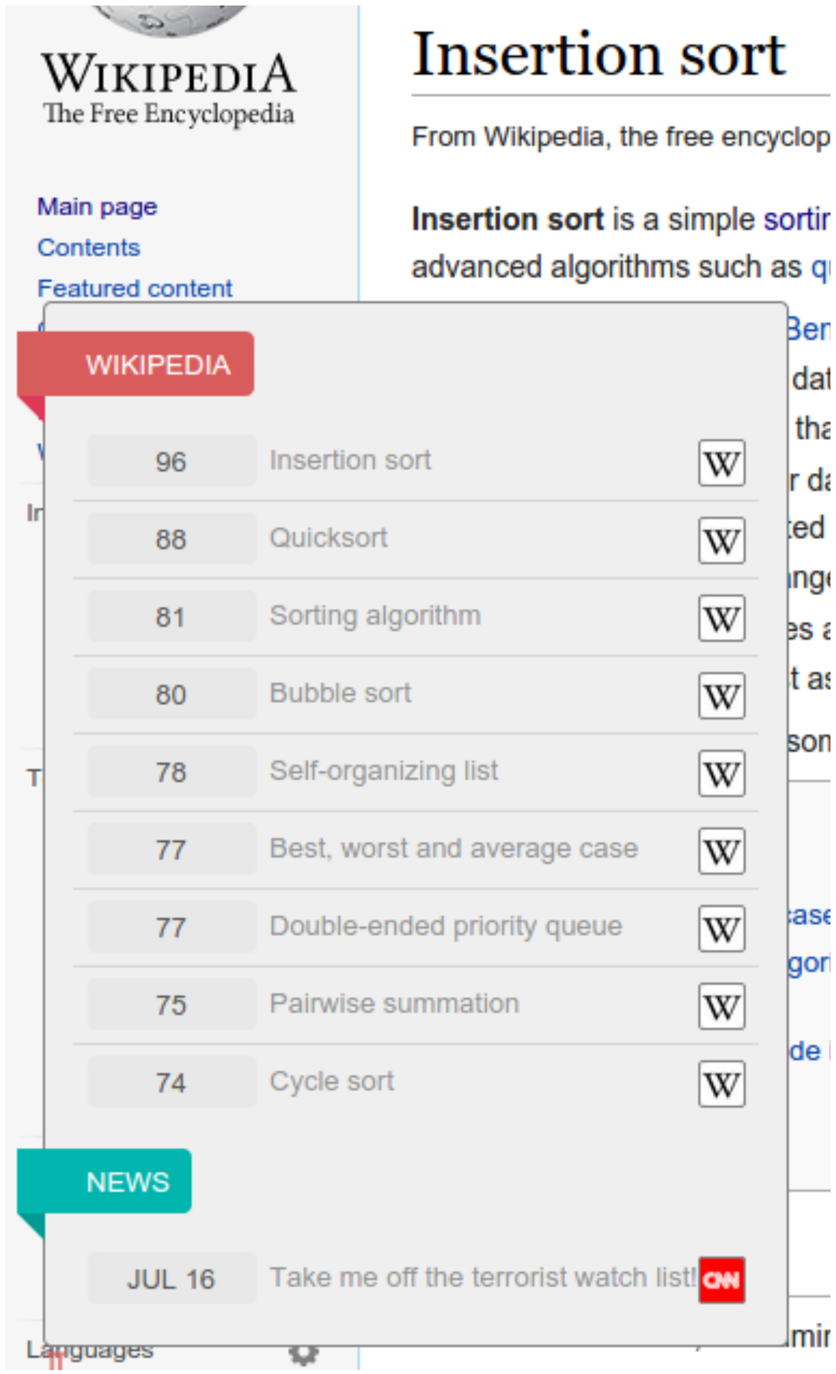

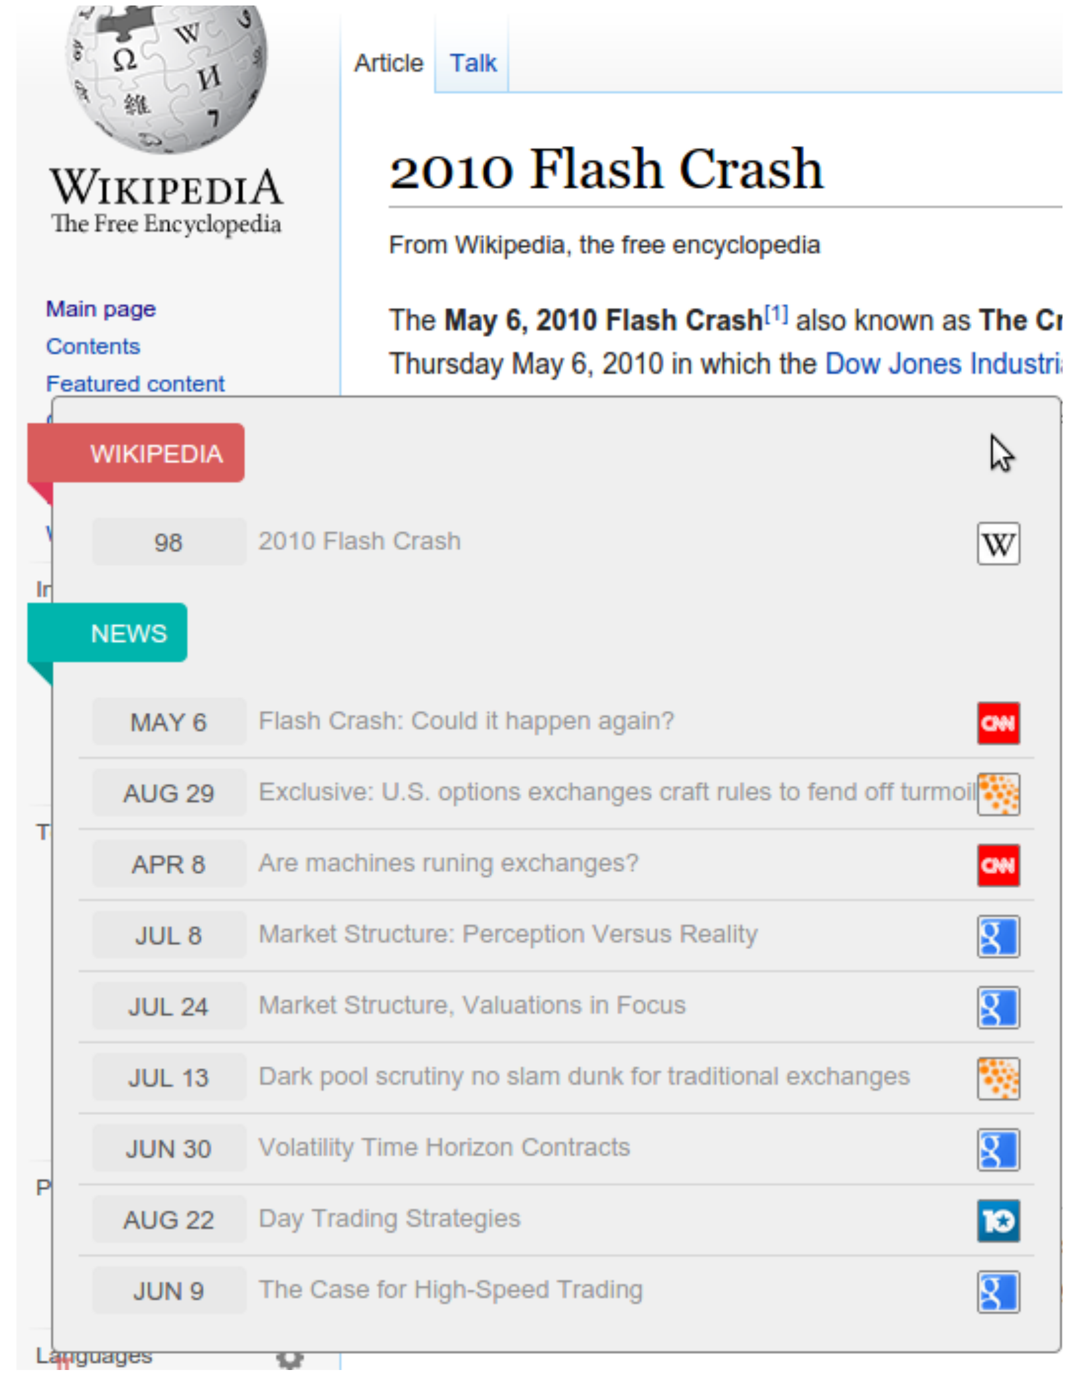

A Graphical Illustration of the Model

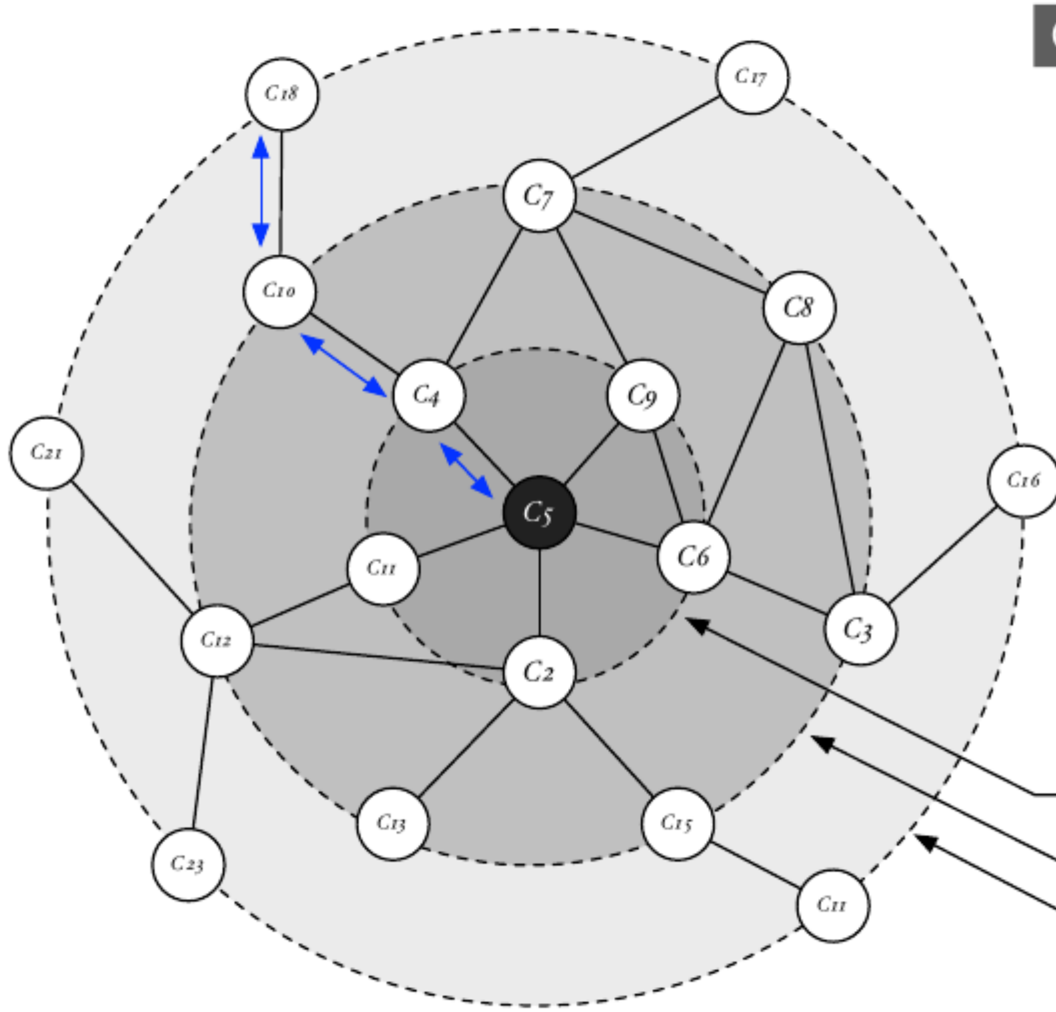

Output Caching strategy (x, y) consists of two parts

1. x : caching decision, what to cache
2. y : retrieving decision, where to retrieve

E.g. $x_{18,3} = 1, y_{5,18,3} = 1$ indicates C_{18} caches O_3, and C_5 retrieves O_3 from O_{18}

Input Topology G, nodes C, content O

N_i^+ defines who has C_i in its neighborhood

$C_5 \in N_6^+, \quad C_5 \in N_9^+, \quad C_5 \in N_{10}^+, \quad C_5 \in N_{11}^+ \quad \ldots\ldots$

N_i defines C_i has whom in its neighborhood

C_5 1-hop neighbors z_1
C_5 2-hop neighbors z_2
C_5 3-hop neighbors z_3
......

C_5 neighborhood N_5

Content localization: fraction of retrieved content