

Effects of Cooperation Policy and Network Topology on Performance of In-Network Caching

Liang Wang, Suzan Bayhan, Jussi Kangasharju

Department of Computer Science, University of Helsinki, Finland

Abstract—In-network caching is a key component in information-centric networking. In this paper we show that there is a tradeoff between two common caching metrics, byte hit rate and footprint reduction, and show that a cooperation policy can adjust this tradeoff. We model the cooperation policy with only two parameters – search radius r and number of copies in the network K . These two parameters represent the range of cooperation and tolerance of duplicates. We show how cooperation policy impacts content distribution, and further illustrate the relation between content popularity and topological properties. Our work leads many implications on how to take advantage of topological properties in in-network caching strategy design.

I. INTRODUCTION

Caching is a key component in information-centric networks (ICN) [1]–[4]. In-network caching not only reduces an ISP’s outgoing traffic, but also reduces traffic within an ISP network. Byte hit rate (BHR) is a common metric for evaluating savings in inter-ISP traffic, however, there is no widely accepted metric for evaluating savings in intra-ISP traffic. Footprint reduction (FPR) [5] has been proposed as one such metric and is the one we use in this paper. We show that BHR by itself is insufficient in capturing the performance of a network of caches; this is often overlooked by existing work. This paper shows that there is a subtle interplay between BHR and FPR and that in some cases these two metrics oppose each other. We argue that a *cooperation policy* among routers can mediate this tradeoff between BHR and FPR and show that two parameters can tune the desired operating region: maximum number of duplicates for each content item (K) and the radius for cooperation (r).

To improve BHR, a cooperation policy covering a large radius enhances network storage utilization by reducing the number of duplicates (*cache diversity* in [6]). However, large-scale cooperation causes communication overheads and increases intra-ISP traffic because requests may be redirected many times. Despite efforts in designing a cooperation policy [7], [8], [10], a proper model for its impact on BHR and FPR is still missing. We characterize cooperation policy by its search strength (r) and capability of reducing duplicates (K). We show how different r and K values lead to different tradeoffs between BHR and FPR, and discuss their implication.

There is considerable interest in exploiting topological properties in cache networks. Initial efforts [11], [12] indicate centrality as a promising metric, but questions like how to measure the topological impact on performance and mechanism of the interplay between topology, and caching strategy still remain open. We use a cache cooperation policy to couple

content with topology and show that this coupling explains how topological properties impact caching performance; the tightness of the coupling indicates degree of topology’s impact.

Our contributions in this paper are as follows:

- 1) We highlight the importance of FPR as a performance metric for in-network caching, and show how BHR and FPR conflict each other at the Pareto frontier.
- 2) We propose a cooperation policy model to show the relationship between cooperation policy, content, and topology. We also categorize different cooperation types.
- 3) We propose a novel way to measure the impact of topology, and perform a thorough numerical analysis to show how it influences system performance.

II. SYSTEM MODEL

Consider a network of M routers, L of which directly receive user requests and are *edge routers*. A router denoted by R_i is equipped with a storage capacity of C_i bytes. We assume N distinct files, denoted by f_i and being s_i bytes in size. All files are stored permanently at the Content Provider (CP) represented as the $(M+1)^{\text{th}}$ router (R_{M+1}). Denote the request probability of f_i by this file’s *popularity* p_i , and denote the popularity vector by $\mathbf{p} = [p_i]$. When a request for f_i arrives to an edge router R_j , R_j first searches f_i in its cache. If R_j possesses it, R_j transmits f_i to the user; this a *hit*. Otherwise, in case of a *miss*, R_j contacts routers in its r -hop neighborhood to see if any of them has f_i ; this is the *cooperation policy*. We call the set of all routers located at most r -hops away from R_j as the *searchable set* of R_j , denoted by \mathcal{S}_j^c . If f_i is stored in \mathcal{S}_j^c , it is retrieved to R_j from the closest router (if multiple routers holding f_i) and forwarded to the user. Let $\mathcal{R}_{j,CP}$ be the set of all routers on the path between a leaf router R_j and the CP (excluding the CP). If no router in \mathcal{S}_j^c stores the item, the request is routed to the next router in $\mathcal{R}_{j,CP}$ and searched there as well as in the new searchable set; there may be overlap between searchable sets of two neighboring routers, depending on r . We define the *reachable set* of a router denoted by \mathcal{S}_j^r as the set of all routers in the searchable sets of routers in $\mathcal{R}_{j,CP}$. If no router in \mathcal{S}_j^r has f_i , it is downloaded from the CP and routed to the user following the backward path.

III. OPTIMAL IN-NETWORK CACHING

A. Cooperation Policy Design

Performance of a cooperation policy is determined by contents in the searchable set which is a function of r . The

diversity of contents cached in this set increases caching efficiency which then calls for a caching scheme avoiding duplicate copies in the set [8], [9]. However, popular content may better be cached in multiple routers to be more accessible from all network edge routers. We model this tradeoff with parameter K which is the maximum number of content replicas in the network. In reality every file would have its own maximum number of copies which emerges automatically if r is fixed; we use a fixed K to illustrate system behavior across the whole parameter range. Using these two parameters, we name a cooperation policy with parameters K and r as (K, r) -Cooperation Policy which can be classified into four as follows:

- 1) Type I, **small** r , **small** K : Weak cooperation due to limited access to other caches and limited availability of popular content; the system is not using all its resources.
- 2) Type II, **small** r , **large** K : This is en-route caching. The most popular content is pushed to the network edge.
- 3) Type III, **large** r , **small** K : Network storage is effectively a single cache. Popular content is in network core.
- 4) Type IV, **large** r , **large** K : Strong cooperation. BHR and FPR cannot be improved at the same time since caching system is fully-utilized and reaches its Pareto frontier.

The complexity of cooperation can be calculated via communication and computation overhead [8]. Initially, all routers exchange their set of stored contents with routers in their searchable set. Assuming that each content is unit size and dropping the router index, this initialization step requires $O(MC|\mathcal{S}^c|)$ messages and results in $O(M^2C)$ message exchanges in the worst case. Upon a change in the cache of a router, this router informs all its r -hop neighbours about the evicted and admitted items. This per change announcement requires $O(|\mathcal{S}^c|)$ message in the worst case. In terms of computation, the cooperation does not involve any processing rather than discovering which of the replicas is closest to a specific router. Therefore, computation overhead is $O(|\mathcal{S}^c|)$.

B. Optimal Caching under (K, r) -Cooperation Policy

Assume a centralized entity deciding which items are stored at each router R_i when a user requests item f_u at time t . This entity knows the content distribution $\mathbf{X}^t = [x_{i,j}^t]$ where $x_{i,j}^t$ is 1 if f_i is stored at node R_j , and zero otherwise. An optimal caching strategy (C_{OPT}) determines whether to cache f_u in the routers between the edge router R_l receiving the request and router R_{hit} storing f_u , and which items to evict in case of full cache occupancy. We refer the set of all these intermediate nodes on the path between R_l and R_{hit} as \mathcal{S} .

C_{OPT} minimizes the total cost of serving user requests by exploiting its knowledge of current content distribution \mathbf{X}^t , file popularities (\mathbf{p}), and file size (s_i) information. Let $c_{j,k}$ denote the cost of downloading one byte at R_j from R_k . An item can be served from edge router R_j or retrieved from another router R_k including the CP. Let our decision variable $x_{i,j,k}^{t+1}$ be 1 if R_j downloads f_i from R_k . The cost function reflects the distance between the two entities and can be calculated using shortest path algorithms. For a (K, r) -Cooperation Policy, as

the routers not in \mathcal{S}_j^r are not reachable from this edge router, we set $c_{j,k} = \infty$ if $R_k \notin \mathcal{S}_j^r$. For harmony of notation, we re-define the content distribution by $\mathbf{X}^t = [x_{i,j}^t]$ (and drop t if we do not refer to a specific time). C_{OPT} is formulated as:

$$\min \left(\sum_{i=1}^N \sum_{j=1}^L \sum_{k=1}^{M+1} s_i p_i c_{j,k} x_{i,j,k}^{t+1} x_{i,k,k}^{t+1} + s_u p_u \sum_{j=1}^L \sum_{\substack{\forall R_k \in \\ \mathcal{S} \cup \mathcal{R}_{M+1}}} c_{j,k} x_{u,j,k}^{t+1} \right) \quad (1)$$

s.t. Cache capacity constraints:

$$\sum_{i=1}^N s_i x_{i,j,j}^t x_{i,j,j}^{t+1} + s_u x_{u,j,j}^{t+1} (1 - x_{u,j,j}^t) \leq C_j, \forall R_j \in \mathcal{S} \quad (2)$$

$$\sum_{i=1}^N s_i x_{i,j,j}^{t+1} \leq C_j, \forall R_j \notin \mathcal{S} \quad (3)$$

$$\text{Maximum replica constraint: } \sum_{j=1}^M x_{i,j,j}^{t+1} \leq K \quad \forall i, \quad (4)$$

$$\text{Feasibility constraints: } x_{i,j,k}^{t+1} \leq x_{i,k,k}^{t+1} \quad \forall i, \forall k \quad (5)$$

$$x_{i,j,j}^{t+1} = x_{i,j,j}^t \quad \forall i, \forall R_j \notin \mathcal{S} \quad (6)$$

$$\text{Service constraint: } 1 \leq \sum_{k=1}^{M+1} x_{i,j,k}^{t+1} \quad \forall i, \forall j, \forall k \in L \quad (7)$$

$$\text{Availability constraint: } x_{i,M+1,M+1}^{t+1} = 1 \quad \forall i. \quad (8)$$

Our objective (1) calculates the cost of serving user requests over all the edge routers and minimizes this cost by favoring the most popular files. Note that if $x_{i,j,j} = 1$, then f_i is stored in R_j . *Cache capacity constraints* in (2) and (3) ensure the total size of items to be stored in a router's cache cannot exceed cache capacity. Only routers in \mathcal{S} can consider putting the requested item f_i into their caches. *Maximum replica constraint* in (4) ensures that an item can have maximum K replicas in the network. Note that by removing this constraint, system can figure out optimal K for each neighborhood automatically. *Feasibility constraint* in (5) reflects f_i being retrievable from R_k only if R_k stores f_i whereas (6) states that contents cached by routers not in \mathcal{S} do not change. *Service constraint* in (7) forces the content to be served from some location (i.e., local cache, another router's cache, or the CP) while *availability constraint* in (8) ensures that all items are available from the CP. Decision variables are binary, i.e., $x_{i,j,k} \in \{0, 1\}$. C_{OPT} is an integer linear programming problem which can be solved with optimization software for small instances of the problem but it requires low-complexity distributed schemes for large scale networks. We leave distributed solutions for future work.

Let $\mathcal{F}_j = \{u_{j,1}, u_{j,2}, u_{j,3}, \dots\}$ be the list of user requests arriving at leaf router R_j where $u_{j,i}$ is the i^{th} request for a file with size $s_{u_{j,i}}$. R_j can retrieve it only from its reachable set \mathcal{S}_j^r which is defined as

$$\mathcal{S}_j^r = \bigcup_{R_k \in \mathcal{R}_{j,CP}} \mathcal{S}_k^c. \quad (9)$$

A request will be counted as hit if at least one of the routers in S_j^r stores it. More formally, we define hit function $\delta_{j,i}$ for request $u_{j,i}$ (assuming $u_{j,i}$ is a request for f_i) as follows:

$$\delta_{j,i} = \begin{cases} 1 & \text{if } \sum_{R_k \in S_j^r} x_{i,k,k} \geq 1 \\ 0 & \text{o/w.} \end{cases}$$

Next, we calculate BHR as follows:

$$BHR = \frac{\sum_{j=1}^L \sum_{\forall u_{j,i} \in \mathcal{F}_j} s_{u_{j,i}} \delta_{j,i}}{\sum_{j=1}^L \sum_{\forall u_{j,i} \in \mathcal{F}_j} s_{u_{j,i}}}. \quad (10)$$

If request $u_{j,i}$ is served from a router that is $h_{j,i}$ hops away from the user and the path from R_j to the CP is H_j hops long, we can compute the FPR as follows:

$$FPR = 1 - \frac{\sum_{j=1}^L \sum_{\forall u_{j,i} \in \mathcal{F}_j} s_{u_{j,i}} h_{j,i}}{\sum_{j=1}^L H_j \sum_{\forall u_{j,i} \in \mathcal{F}_j} s_{u_{j,i}}}. \quad (11)$$

IV. NUMERICAL ANALYSIS

A. Setup & Metrics

We performed numerical evaluation on realistic and synthetic topologies. Realistic topologies are from [13], and synthetic topologies are scale-free networks of 50 nodes. Each node can store 25 objects. We present results on synthetic networks; realistic topologies produce similar results. Content popularity is modeled according to [14], and content set contains 5000 objects. We calculate the betweenness centrality (C_B) of each router in order to analyze its impact on cached content in a specific router under various (K, r) pairs. We define *coupling factor* (CPF) as the Pearson correlation between C_B and average popularity per bit in a node's cache; it measures topological impact on system performance. The rationale is that optimal system performance is achieved by placing content at specific locations in a network according to its popularity and that C_B is a good metric to characterize a node's position in a graph. Strong correlation between the two indicates that content is tightly "coupled" with topology and topological properties influence system performance.

In the simulations, 30% of the edge routers are randomly selected and connected with client, and the server randomly connects to one of the 5 core nodes with highest C_B . Experiments were repeated at least 50 times.

B. Pareto Frontier

Fig. 1 shows how K and r impact caching performance. The solution to C_{OPT} provides the optimal cache profiles for given K and r (e.g., point A in Fig. 1), but it does not indicate the best values for these two parameters, i.e., we can improve performance by tuning K and r , because the system may be underutilized. However, our optimization model can be used to find Pareto frontier of the performance (green arc BC in Fig. 1). When we reach the Pareto frontier, we cannot improve BHR or FPR without hurting the other. The fan-shaped area defined by ABC is the area which a cooperation policy can explore to find the best tradeoff between K and r . Point D where we eventually reach the Pareto frontier depends on how

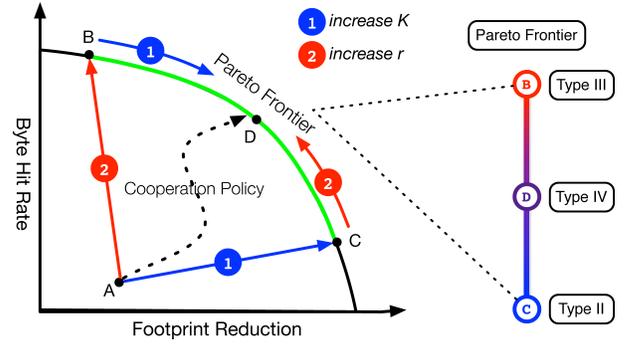


Fig. 1: Conflicting BHR and FPR at the Pareto frontier. Type of cooperation at different points shown on right.

cooperation policy balances BHR and FPR. Lines AB and AC are not parallel to the x- and y-axis, since changing either of r or K affects both BHR and FPR, as we show below.

The upper graph in Fig. 2a shows how BHR and FPR vary as we move along the segments AB , BC , and CA , by varying r and K . Starting from A and moving clockwise (left to right in the figure), we increase the search radius which improves BHR, but decreases FPR due to additional search traffic or letting content be cached at routers with higher $h_{j,i}$ in (11). From B to C , along the Pareto frontier, we observe the tradeoff between BHR and FPR, with FPR reaching its maximum at C . From C to A , r is 0 so the system reduces to en-route caching where larger number of copies (near C) is beneficial, hence as we move towards A , both BHR and FPR decrease. Fig. 2b shows heatmaps of CPF, BHR, and FPR as function of K and r . Lighter values indicate higher values. It shows how BHR and FPR conflict each other, i.e., one achieving the highest performance while the other has the worst, in regions corresponding to the Pareto frontier.

C. Coupling Content and Topology

The lower plot in Fig. 2a shows how CPF evolves along the same path. Values close to -1 or 1 indicate strong dependence between popularity and betweenness. A router with a high C_B is in or close to the core of the network whereas a router with low C_B is close to the network edge. At B , where CPF is close to 1, popular content is in nodes with high C_B , i.e., the core, whereas at C , where CPF is close to -1, it is at the edge where C_B is low. Along the Pareto frontier BC , we observe a "migration" of content from core to edge. At D where CPF is 0, both BHR and FPR are close to halfway point between their respective minima and maxima at B and C . We have observed this phenomenon across a wide range of experimental settings, but its full investigation is left for further study.

Fig. 3 shows cooperation policy's impact on content placement along the Pareto frontier BC . Point B (Fig. 3a), representing Type III cooperation favors BHR and places content in the core. Point D along BC (Fig. 3b) strikes a tradeoff between BHR and FPR and the content is neither in the core nor on the edge; this is Type IV cooperation. Finally, point C (Fig. 3c) favors FPR and pushes popular content to the edge.

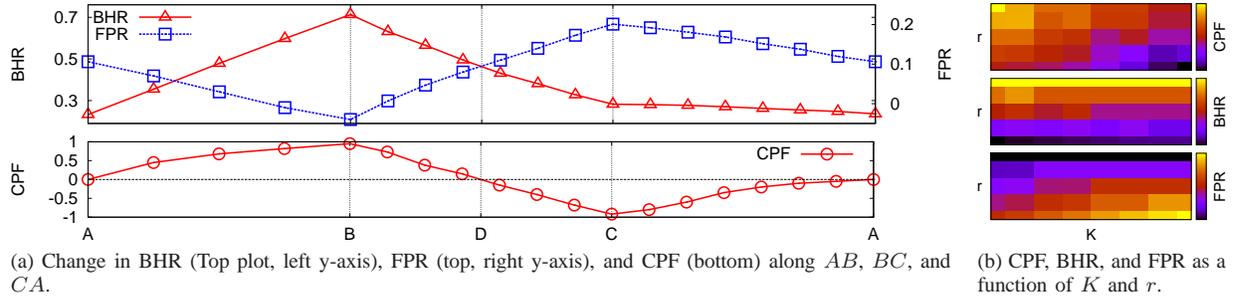


Fig. 2: Performance of (K, r) -Cooperation along the boundary defined by ABC (a), and for all (K, r) pairs (b).

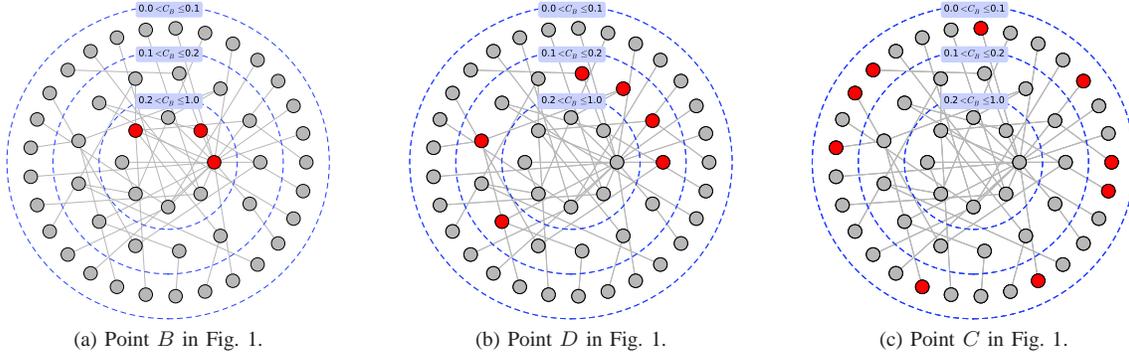


Fig. 3: Content placement at points B , D , and C . Red color (dark dots) marks the most popular content. Nodes are grouped in circles according to C_B . Content migrates from core to edge as we move from B to C .

D. Implications

Our results have profound implications on relationship between cooperation policies, content popularity, and network topology. They also give us hints on when and how topological properties should be taken into account in caching strategy design. We summarize the main implications as follows:

- 1) Cooperation policy pushes performance to Pareto frontier and couples content popularity and topological properties together. How and where it falls on the frontier depends on how it balances BHR and FPR.
- 2) Content popularity and topology strongly correlate with each other only close to the Pareto frontier. Whether the correlation is positive or negative depends on how the cooperation policy favors one of the two metrics.
- 3) The optimization model implies that C_B has more influence on performance when we get closer to points A or B in Fig. 1; only Type II and III cooperation policies can fully utilize C_B to enhance performance.
- 4) We only present results on betweenness centrality. We also experimented with other centrality measures and got similar results. The impact of other properties like diameter or clustering coefficient needs further study.
- 5) We conjecture that tight coupling between content popularity and topology comes similar mathematical structures as both exhibit power-law properties. Were popularity closer to uniform or topology closer to a random network, this tight coupling might disappear. However, this matter requires further study.

V. CONCLUSION

We modeled cache cooperation by its search radius and tolerance of duplicates. We performed a thorough numerical analysis and showed that cooperation policy pushes system performance to its Pareto frontier, and how it couples content with topology. We proposed a way to measure impact of topology on system performance. We show when and how topological information should be taken into account in in-network caching strategy design.

REFERENCES

- [1] V. Jacobson, et al., "Networking named content," in *Proceedings of ACM CoNEXT*, 2009.
- [2] T. Koponen, et al., "A data-oriented (and beyond) network architecture," *Proceedings of ACM SIGCOMM*, 2007.
- [3] Publish/Subscribe Internet Routing Paradigm, "Conceptual architecture of PSIRP including subcomponent descriptions. Deliverable d2.2, PSIRP project," , August 2008.
- [4] C. Dannowitz, "Netinf: An information-centric design for the future internet," in *GI/ITG KuVS Workshop on The Future Internet*, 2009.
- [5] A. Anand, V. Sekar, and A. Akella, "SmartRE: an architecture for coordinated network-wide redundancy elimination," in *Proceedings of ACM SIGCOMM*, 2009.
- [6] G. Rossini and D. Rossi, "Evaluating CCN multi-path interest forwarding strategies," *Computer Communications*, v. 36, n. 7, pp. 771–778, 2013.
- [7] Y. Li, H. Xie, Y. Wen, and Z.-L. Zhang, "Coordinating in-network caching in content-centric networks: Model and analysis," in *IEEE ICDCS*, 2013.
- [8] V. Sourlas, L. Gkatzikis, P. Flegkas, and L. Tassiulas, "Distributed cache management in information-centric networks," *IEEE Transactions on Network and Service Management*, v. 10, no. 3, pp. 286–299, 2013.

- [9] J. M. Wang, J. Zhang, B. Bensaou, "Intra-AS cooperative caching for content-centric networks," in *ACM SIGCOMM Workshop on Information-Centric Networking*, 2013.
- [10] S. Saha, A. Lukyanenko, and A. Yla-Jaaski, "Cooperative caching through routing control in information-centric networks," in *IEEE INFOCOM*, 2013.
- [11] D. Rossi and G. Rossini, "On sizing CCN content stores by exploiting topological information," in *Infocom Computer Communications Workshops*, 2012.
- [12] W. K. Chai, D. He, I. Psaras, and G. Pavlou, "Cache "less for more" in information-centric networks," in *IFIP Networking Conference*, 2012.
- [13] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in *Proceedings of ACM SIGCOMM*, 2002.
- [14] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon, "I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system," in *Proceedings of ACM IMC*, 2007.