

# Exploiting Communities for Enhancing Lookup Performance in Structured P2P Systems

H. M. N. Dilum Bandara and Anura P. Jayasumana

Colorado State University

[Anura.Jayasumana@ColoState.edu](mailto:Anura.Jayasumana@ColoState.edu)



University of  
Massachusetts Amherst



University of Oklahoma



Colorado State University



University of  
Puerto Rico Mayaguez

CASA is primarily supported by the Engineering Research Centers Program  
of the National Science Foundation under NSF award number 0313747.



# Contribution

## Community-aware caching scheme to enhance lookup performance in structured P2P systems

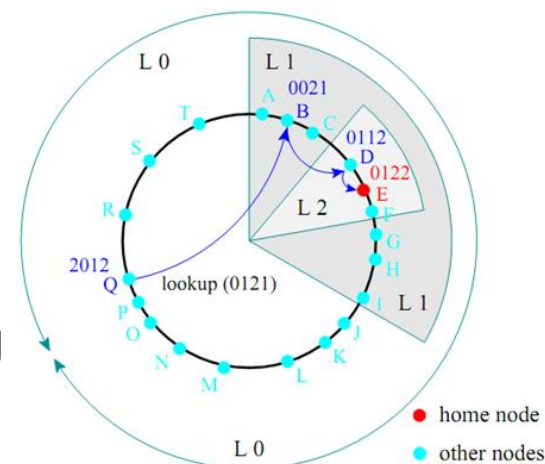

1. Build sub-overlays among community members while preserving overlay properties
  2. Weighted least frequently used caching based on local statistics
- Enhances both communitywide (23-51%) & system-wide lookup (40%) performance
  - Works with structured P2P systems that provide alternative paths to a given destination
  - Works with any skewed popularity distribution
  - Adaptive to changing popularity
  - Need small caches

# Motivation

- Many small communities are emerging within P2P systems
- Community – subset of peers that share some similarity
  - Semantic
    - Many BitTorrent communities – music, movies, games, Linux distributions, private communities
  - Geography
    - For 60% of files shared by eDonkey peers, more than 80% of their replicas were located in a single country [Handurukande, 2006]
  - Organizational
    - Peers within an AS, members of a professional organization, group of universities
    - To share resources & limit unrelated external traffic

# Motivation (cont.)

- Content popularity in P2P follows Zipf's-like distribution
- Improve lookup
  - Restructure overlay based on similarity
  - Cache most globally popular content
- However
  1. Communities are not isolated
  2. Individual communities don't rank high in popularity
  3. Not every node can or interested in caching



[Ramasubramanian, 2004]

# Content Popularity in Communities

## 1. Communities are not isolated

faltu fringe vampire diaries hall pass no  
Angeles Limitless criminal minds rango  
biutiful bdsm thor big bang theory csi drive ar  
lawyer archer Megamind fast five x-art xxx 127  
little fockers your highness Stargate Universe hop  
with it fxx windows 7 hindi ita one tree hil  
justified wwe unknown insidious the green I  
rio MAXSPEED britney spears modern family teen  
how i met your mother chuck paul d  
family guy bruno mars black eyed peas hentai the social n  
idol toy story 3 french Beastly The tourist noir

BitTorrent  
Communities

Community*	EX	FE	SP	TB	TS	TE	TR
fenopy.com (FE)	0.38				EX – extratorrent.com		
seedpeer.com (SP)	0.00	0.00					
torrentbit.net (TB)	0.40	0.29	0.00				
torrentscan.com (TS)	0.48	0.33	0.00	0.48			
torrentsection.com (TE)	0.53	0.23	0.00	0.31	0.25		
torrentreactor.net (TR)	0.10	0.08	0.00	0.06	0.09	0.06	
youbittorrent.com (YB)	0.36	0.35	0.00	0.29	0.42	0.20	0.04


# Content Popularity in Communities (cont.)

## 2. Communities have different Zipf's parameters



- $\alpha = 0.53, 0.66, 0.79, 0.98$
- Aggregation of multiple Zipf's distributions is not necessarily Zipf
- Caching on a structured P2P system with alternative paths [Rao, 2007]

$$f_r = \frac{1/r^\alpha}{\sum_{n=1}^N 1/n^\alpha}$$


$$H = \log N - \sum_{r=1}^C f_r \log f_r - \log_k L$$





+



# Structured Overlay – Chord DHT




# Sub-Overlay Formation



- Goal – not to isolate communities or mix contents
- Each community forms a sub-overlay
  - Form links/fingers to community members
- Enable nodes to identify what's popular in their community & cache accordingly
  - Forward queries to community members hoping that they may have already cached required contents



# Sub-Overlay Formation (cont.)



$B \rightarrow D \rightarrow F = 2$  hops

$B \rightarrow E \rightarrow F = 2$  hops

If E cache F's content

$B \rightarrow E = 1$  hop

No of distinct node found by  
probing  $i$ -th finger & it's successor

$$2(i + 2 \log N - m) - 1$$

$N$  – No of nodes

$m$  – Key length

$$1 \leq i \leq m$$

- Nodes have 1 or more community IDs
  - Communities based on different similarity measures – semantic, geography
  - Support exceptions – user in USA can be a member of a community in India
- Identify community members that are at an exponentially increasing distances in key space
  - Sample nodes pointed by links & their successors
  - Long distant links (large  $i$ ) are more important & easy to find

# Caching Algorithm

- Cache based on community interest
  - Queries go through community members → Nodes get to know what's popular in their community
- Local statistics are sufficient to estimate relative popularity
  - Focus on community interest
  - No assumption on popularity distribution
- Weighted least frequently used caching
  - Evaluate demand at arrival of each query  $q$  → Adaptive
  - Weight  $\alpha$  determine bias towards short or long term trends

$$\begin{cases} demand_i^k = (1 + \alpha) \times demand_{i-1}^k & \text{If } q \text{ is for } k \\ demand_i^k = (1 - \alpha) \times demand_{i-1}^k & \text{else} \end{cases} \quad 0 \leq \alpha \leq 1$$

- If  $demand^k > D_{cache}$  – Indicate node's interest to cache by append to query  $q$
- Query response is send to query originator & all nodes that want a copy to cache

# Caching Algorithm (cont.)

- Reevaluates what keys to cache at arrival of a query
  - Naturally adapts to varying trends of community interests
  - Computationally efficient
- Track contents even if not cached
  - Threshold to remove least popular ones
- $D_{cache}$  – Caching threshold
  - Prevents cache thrashing
  - $D_{cache} > \alpha$

```
void forward(key, msg, nextHop*)
```

```

1  If msg.type = PUT                //put message
2  return
3  If msg.type = GET                //get message
4  addLookup(key)                  //Track demand
5  If key ∈ C                      //In cache
6  sendDirect(msg.source, key, C[key])
7  For each i in msg.cList[ ]      //Send to each cache requester
8  sendDirect(msg.cList[i], key, C[key])
9  nextHop ← NULL                  //Drop original get message
10 Else                           //Not in cache
11  If C.size( ) = Cmax           //Cache already full
12  key_lowest ← getCachedKeyWithLowestDemand(L[ ])
13  If L[key] > L[key_lowest]       //Higher demand
14  msg.cList[ ] ← myNodeID        //Request a copy
15  C[key_lowest].remove           //Remove lowest key
16 Else
17  If L[key] > Dcache             // Higher demand
18  msg.cList[ ] ← myNodeID        //Request a copy

```

```
void addLookup(key)
```

```

19 For each i in L[ ]
20 If i = key                      //Increase demand for key
21 L[i] = (1 + α) × L[i]
22 Else                            //Decrease demand for others
23 L[i] = (1 - α) × L[i]
24 If L[i] < Dremove             //Very low demand
25 L[i].remove                     //Remove key

```

# Simulation Setup

Community	C <sub>1</sub>	C <sub>2</sub>	C <sub>3</sub>	C <sub>4</sub>	C <sub>5</sub>	C <sub>6</sub>	C <sub>7</sub>	C <sub>8</sub>	C <sub>9</sub>	C <sub>10</sub>
No of nodes (apx.)	600	600	600	1,200	1,200	1,200	1,200	1,200	2,400	4,800
Zipf's parameter	0.85	0.95	1.10	0.5	0.80	0.80	1.0	0.90	0.90	0.75
No of distinct keys	40,000	30,000	30,000	40,000	40,000	40,000	50,000	50,000	50,000	50,000
Similarity with community (x)	0.2 (C <sub>8</sub> )	0	0.1 (C <sub>7</sub> )	0.2 (C <sub>9</sub> )	0.3 (C <sub>8</sub> ) 0.5 (C <sub>7</sub> )	0	0.1 (C <sub>3</sub> ) 0.5 (C <sub>5</sub> )	0.3 (C <sub>5</sub> ) 0.2 (C <sub>1</sub> )	0.4 (C <sub>1</sub> ) 0.2 (C <sub>4</sub> ) 0.3 (C <sub>10</sub> )	0.3 (C <sub>9</sub> )
Queries for rank 1 key	4,516	8,535	17,100	603	6,454	6,454	21,059	11,956	23,911	17,030

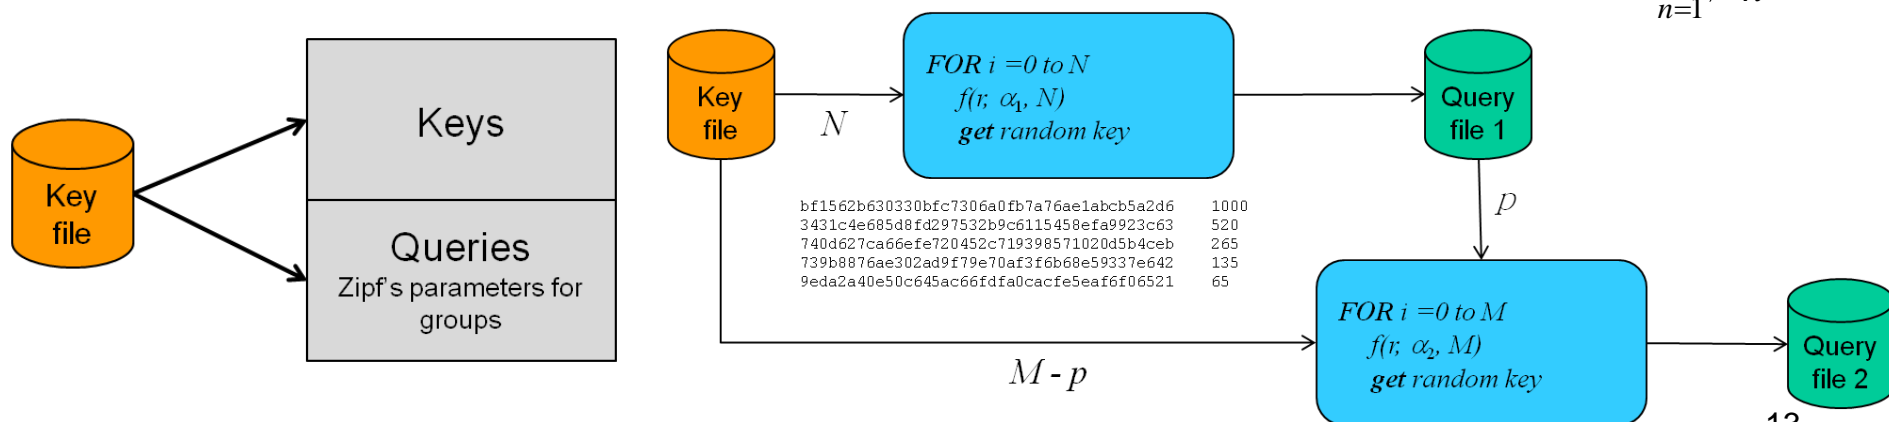
- OverSim P2P simulation environment
- Sub-overlay formation & caching implemented on top of Chord overlay
- 15,000 nodes
- 10 communities of different sizes
- Different Zipf's parameters
- Queries after system got stabilized – around 2000 sec
- 10 samples

# Community, Keys & Query Generation

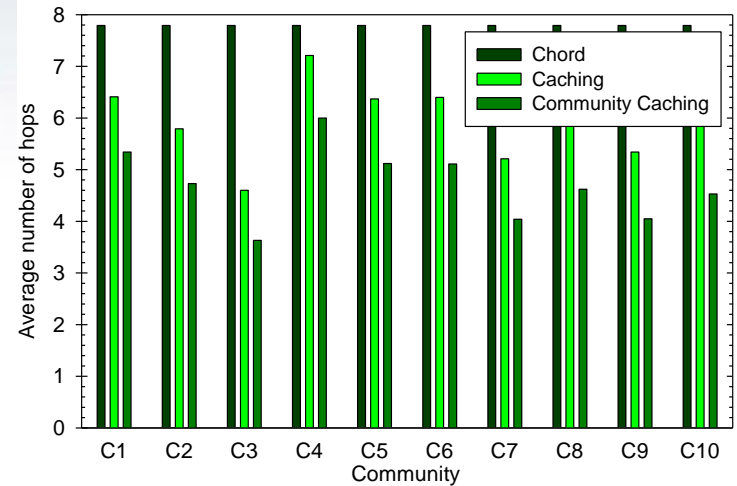
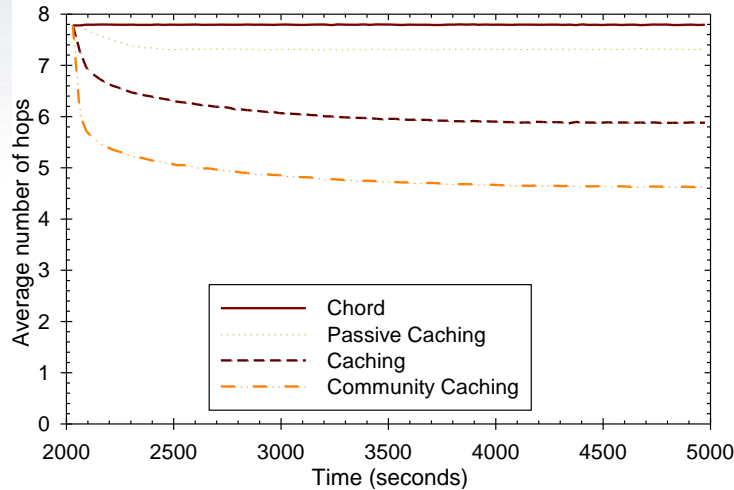
- Peers know their group ID at initialization
- Each peer
  - Maintain a key index – no capacity limit
  - Maintain a cache – fixed capacity
- Generate fixed set of keys a-priori
  - Peers read keys from a file & store in appropriate nodes
- Queries

- Use set of Zipf's parameters observed form BitTorrent

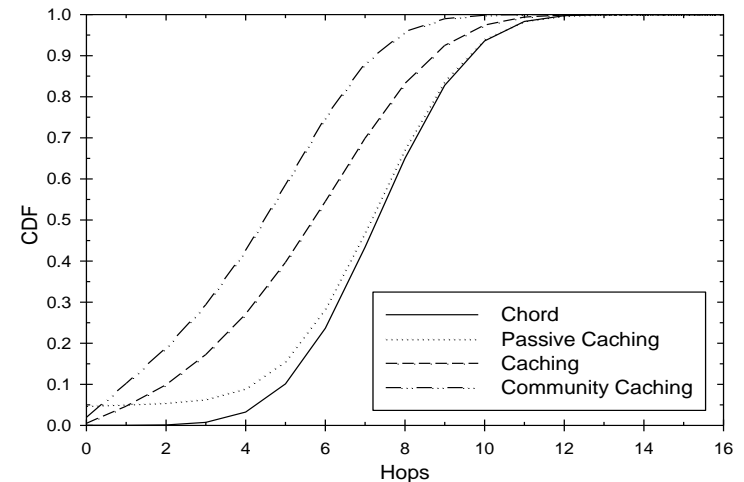
$$f(r, \alpha, N) = \frac{1/r^\alpha}{\sum_{n=1}^N 1/n^\alpha}$$



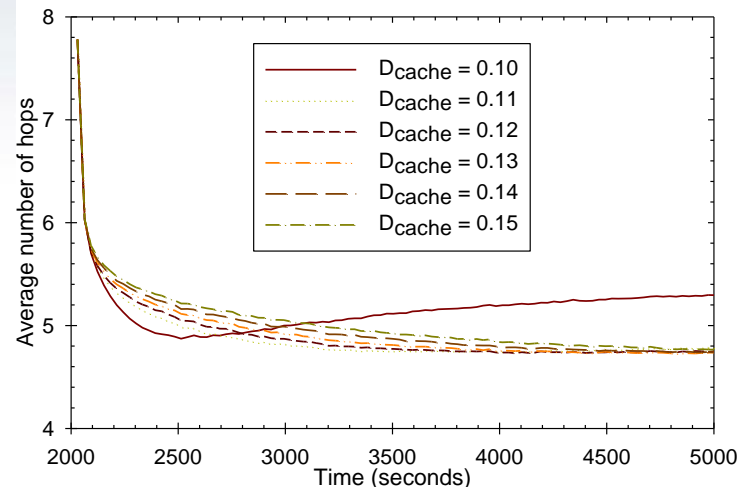
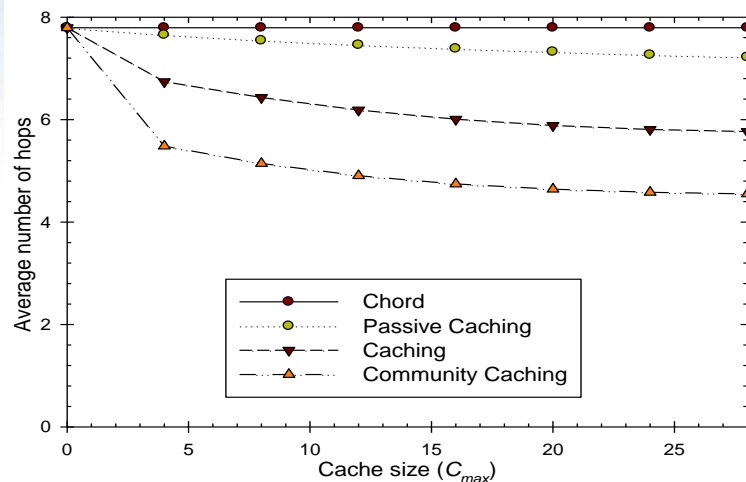
# Performance Analysis



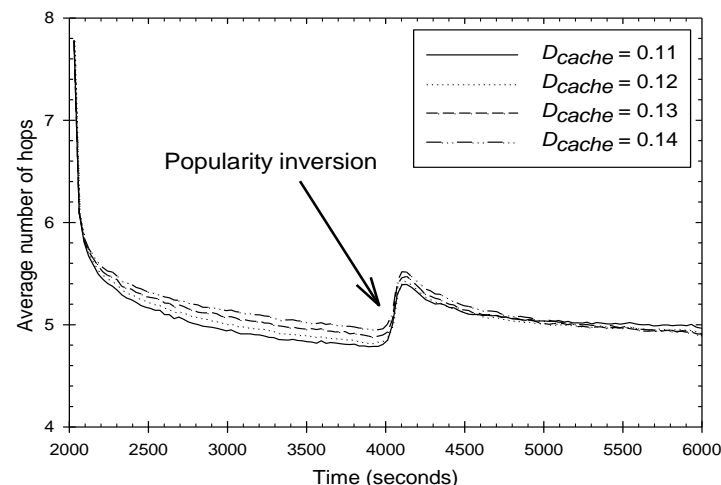
- Reduced path length
  - Overall system – 40.5%
  - More popular communities – 48-53%
  - Least popular community – 23% reduction (7% with caching)
- Performance depends on skewness
  - $C_1, C_5, \text{ \& } C_6$
- Most queries are responded within few hops



# Performance Analysis (cont.)



- Small cache size per node
- $D_{cache}$  reduce cache thrashing, overhead, & long-term path length
- Rapidly respond to popularity changes
- Better load distribution
  - Max with Chord – 27,574
  - Max with Community Caching – 1,677



# Summary

- Community-aware caching solution for structured P2P
  - Allows queries to be forwarded through community members
  - Enable nodes to cache resources that of interest to their community
- Properties
  - Improve both communitywide & system-wide performance
  - Works with any structured P2P system that provides alternative paths to a given destination
    - Preserve overlay bound  $O(\log N)$
  - Independent of popularity distribution & how communities are formed
  - Based on local statistics
  - Adaptive
  - Introduces minimum cache storage, network, & computational overhead
- Current/future work
  - Analyze performance under peer churn, heterogeneous caches, & geography based communities
  - In-network community identification & formation



# *Questions ?*

Anura.Jayasumana@ColoState.edu  
[www.cnrl.colostate.edu](http://www.cnrl.colostate.edu)