

<http://ansinet.com/itj>

ITJ

ISSN 1812-5638

INFORMATION TECHNOLOGY JOURNAL

ANSI*net*

Asian Network for Scientific Information
308 Lasani Town, Sargodha Road, Faisalabad - Pakistan

A Hybrid P2P-CDN Based on Locality-awareness and Interest-awareness

¹Shen Fuke and ²Zhu Yuhong

¹Network Information Center, East China Normal University, Shanghai, China

²Information Center, Shanghai Municipal Education Commission, Shanghai, China

Abstract: Non-profit websites (e.g., education or social organizations, etc.) can not afford the expensive expenses of deploying and administrating a dedicated Content Distribution Networks (CDN) infrastructure. What these websites need is a distributed content distribution infrastructure that can deliver the content at large scale without the large costs. P2P-CDN which combines P2P with CDN technologies is an appropriate solution. In this study, We propose a hybrid P2P-CDN architecture named Aroma-CDN which consisting of a Distributed Sloppy Hash Table (DSHT) structured overlay and unstructured overlays (petal) to provide locality-aware query routing. We can get fast lookup speed and good scalability while avoiding hot-spot congestion and tree saturation. In Aroma-CDN, When a new client requests some content from a website, a locality-aware multiple DSHT infrastructure will quickly find a petal that has the content available. Within a petal which represents a cluster of peers that have the same interest and reside close to each other, peers use gossip protocols to exchange information about their content and contacts rapidly. Simulation-based evaluation shows that Aroma-CDN yields a high performance gains under large population.

Key words: Peer to Peer (P2P), overlay, content distribution networks, distributed sloppy hash table, locality-awareness, interest-awareness

INTRODUCTION

Dedicated Content Distribution Networks (CDN) are expensive technologies for highly popular websites' file-sharing and web-distributing to large audiences. Unfortunately, non-profit websites (e.g., related to education organizations, social organizations, etc.) can't afford the expensive expenses of deploying and maintaining a dedicated CDN infrastructure (El Dick *et al.*, 2011; Stading *et al.*, 2002). Their under-provisioned servers easily become overloaded with large amount of queries. What these websites need is a distributed content distribution infrastructure that can deliver the content at large scale with a limited cost. P2P-CDN which combines P2P with CDN technologies is an appropriate solution.

In this study, we propose a hybrid P2P-CDN architecture named Aroma-CDN which consisting of a Distributed Sloppy Hash Table (Freedman *et al.*, 2004) (DSHT) structured overlay and unstructured overlays (petal) to provide locality-aware query routing with limited background traffic cost. The Aroma-CDN is based on locality-awareness and interest-awareness to improve speed and efficiency of peers searching resource copy and content transfer.

More precisely, this study makes the following contributions:

- We propose a novel scalable hybrid architecture Aroma-CDN which combines hierarchical DSHT structured overlay and unstructured overlays
- First, Aroma-CDN allows nodes to locate nearby cached copies of content objects without querying more distant nodes
- Second, Aroma-CDN prevents hot spots in the directory peers infrastructure

The unstructured overlay clusters (called petals) share the same locality and are interested in the same web objects. Within a petal, peers use gossip protocols to exchange information about their content and contacts, allowing Aroma-CDN to maintain accurate information despite dynamic changes in order to support eventual queries.

Simulation evaluation shows Aroma-CDN's high performance gains and fast convergence.

RELATED WORK

Many efforts have focused on exploiting locality in order to improve the performance of P2P CDN. The centralized P2P-CDN (Ryu and Yang, 2005) relies on the web-server to centrally manage the directory of peers in

which the objects have been cached. These approaches lack robustness and scalability. Stavrou *et al.* (2002) and Wang *et al.* (2002) use unstructured overlays for their flexibility, robustness and load balancing without locality-awareness, while induces heavy traffic and high latency with large scale audience. Structured approaches (Iyer *et al.*, 2002; Rao *et al.*, 2007; Stading *et al.*, 2002) try to rely on DHT to achieve fast lookup without any locality or interest considerations. If the directory information is abruptly lost at the failure of its storing peer, performance may be severely degraded. Additionally, the whole DHT and random physical location navigating implies more routing load and higher response times. Rao *et al.* (2007) is P2P overlay routing infrastructure based on XOR metric topology, it is efficient routing algorithm but it did not take account to the network locality. As one of currently available for public used P2P CDN, CoralCDN (Freedman *et al.*, 2004) relies on a hierarchy of tree-based overlays that cluster nearby nodes. Object searching starts at the highest level (Lower RTT) overlay of the proxy to benefit from network locality then progresses down the hierarchy. But in Freedman *et al.* (2004), users are not involved in the P2P network, an increase of the number of users requires more investment in terms of proxy caches. Flower-CDN (El Dick *et al.*, 2011) combines structured D-ring directory nodes with unstructured gossiping content nodes, to provide efficiency and petal wide locality-aware routing for user queries. But lack of locality-awareness in directory nodes can result in slow lookup. D-ring's frequent references to the same directory can generate tree saturation.

Our work is to combine Flower-CDN with Coral CDN, which using the Coral CDN's idea among the directory peers while using the Flower-CDN petal idea among content peers. We use Kademia (Rao *et al.*, 2007) XOR metric topology to divided directory overlay into four independent routing infrastructure in present study .

DESIGN AND IMPLEMENTATION

Now we introduce the main idea of the new hybrid architecture. Aroma-CDN is designed to support a set O of content objects (e.g., set of web-pages or documents other than websites). If there is enough number of clients willing to participate as cache of these content objects, these content objects will be supported by Aroma-CDN. Also the sum of the clients has to be active enough to actually achieve an acceptable hit ratio. We use the binning technique (Stading *et al.*, 2002) as our locality-awareness. A peer measures its RTT to a set of well-known landmarks and orders them by increasing latency. Physically close peers are likely to have the same

landmark ordering. Thus, each possible ordering identifies a locality loc : $1 \leq loc \leq k$ with k the total number of localities. Figure 1 illustrates the architecture of Aroma-CDN. Peers belonging to the same locality loc and interested in the same content object β build an unstructured overlay noted petal (β, loc) , using gossip protocols. These peers, called content peers and noted $c(\beta, loc)$, cache, manage and exchange content of β , thus relieving the source server's query load. Aroma-CDN chooses one peer of each petal (β, loc) , the role of a directory peer (noted DP or $d(\beta, loc)$): $d(\beta, loc)$ knows about all content peers $c(\beta, loc)$ and keeps index information about their stored content. Directory peers are organized in so-called hierarchical DSHT, A DSHT is a structured overlay based on distributed hash table. For locality-aware routing, we use 4 levels of DSHTs called clusters. A set of nodes should form a cluster if their average, pair-wise RTTs are below some threshold. Our four-level hierarchy has thresholds of ∞ , 200, 100 and 50 msec for level-0, -1, -2 and -3 clusters, respectively.

With the help of Aroma-CDN, when a client looks for an object β with a given key, from the customer's point of view, regardless of the website server anywhere in the network, clients are able to quickly locate resource copy which lies the nearest to the client on physical network, thus, improve the access speed. To the web site's standpoint, the web content should be pushed to the network quickly and balanced.

Key management: In order to ensure a fast lookup and to prevent tree saturation, DPs are constructed into hierarchical DSHT. For each popular web object $\beta \in O$, the directory overlay enables k (where k is the number of localities) participant peers to join as directory peers for β : each locality loc is covered by a directory peer $d(\beta, loc)$, for locality-aware redirection of queries. In the example of Fig. 1, Aroma-CDN covers three content objects (denoted by Red, Green and Grey) and seven localities, i.e., $k = 7$. Thus, each content object has 7 directory peers.

In DHT-based systems, peer identifiers (noted ID) are chosen from an identifier space $S = [1..2^{m-1}]$, where, m is the ID length in bits. Based on these identifiers data value is then typically determined by a hash function which maps data identifiers to peer identifiers. Every object receives a key and the peer with the ID closest to the object key is responsible for storing the pointers to the locations of object replicas. In Aroma-CDN, we want that a query for content object $\hat{\alpha}$ posed by a peer in locality loc quickly finds the DP $d_{\beta, loc}$.

To achieve this, we only have to assign a DP a very specific peer ID, namely an identifier based on the content




Fig. 1: Aroma-CND architecture




Fig. 2: Aroma-CDN's KEY structure

object β and locality it represents. As shown in Fig. 2, the m bits of a peer ID are split into two segments, a locality ID and a content object ID:

- **Locality ID:** Identifier of the locality to which the DP belongs. It is expressed using the highest bit-segment of length m_1
- **Content object ID:** Identifier of the content object β which the DP serves. It is expressed using the lowest bit-segment of length $m_2 = (m - m_1)$. The identifier is obtained by hashing the url of content object (noted hash (β))

As internet object accessing obeys the zipf distribution, Aroma-CDN operates with content object o , other than website ws as Flower-CDN does.

Hierarchical directory peer indexing: In Aroma-CDN, DPs use a sloppy storage technique DSHT that caches key/value pairs at DP whose ID is close to the key being referenced. They frequently satisfy put and get requests at DPs other than those closest to the key. This characteristic differs from DHTs, whose put operations always proceed to nodes closest to the key. So, DSHT cached values reduce hot-spot congestion and tree saturation (Freedman *et al.*, 2004). The Insertion algorithm also performs a two-phase operation to insert a key/value pair (Freedman *et al.*, 2004). In the first or forward phase, it routes to nodes that are successively closer to the key. However, to avoid tree saturation, an insertion operation may terminate prior to locating the closest node to the key, in which case the key/value pair will be stored at a more distant node (Freedman *et al.*, 2004). More specifically, the forward phase terminates whenever the storing node happens upon another node that is either full or loaded for the key:

- A node is full: When it stores l values for k whose TTLs are all at least one half of the new value
- A node is loaded: When it has received more than the maximum leakage rate γ requests for k within the past minute

In our experiments, there is a little bit different with Freedman *et al.* (2004), $l = 8$ and $\gamma = 10$, meaning that under high load, a node claims to be loaded for all but one store attempt every 6 sec. This prevents excessive numbers of requests from hitting the key's closest nodes, yet still allows enough requests to propagate to keep values at these nodes fresh.

Figure 3 illustrates DP's hierarchical routing operations. Each DP has the same node ID in all clusters to which it belongs; we can view a node as projecting its presence to the same location in each of its clusters. This structure must be reflected in Aroma's basic routing infrastructure, in particular to support switching between a node's distinct DSHTs midway through a lookup. A requesting node R specifies the starting and stopping levels at which Aroma should search. By default, it initiates the get query on its highest (level-3) cluster to try to take advantage of network locality. If routing RPCs on this cluster hit some node storing the key k (RPC 1), the lookup halts and returns the corresponding stored value(s) without ever searching lower-level clusters. If a key is not found, the lookup will reach k 's closest node C_3 in this cluster (RPC 2), signifying failure at this level. So, node R continues the search in its level-2 cluster. C_3 likely exists at the identical location in the identifier space in all clusters. R begins searching onward from C_3 in its level-2 cluster (RPC 3), having already traversed the ID-space up to C_3 's prefix. Even if the search eventually switches to the global cluster (RPC 4), the total number of RPCs required is about the same as a single-level lookup service, as a lookup continues from the point at which it left off in the identifier space of the previous cluster. Thus (1) all lookups at the beginning are fast (2) all pointers in higher level clusters reference data within that local cluster.

Content searching process: A DP's content searching process query(o) which to request a content object o , begins with a new client's query delivery.

In case receiving a query(o), $d(\gamma, loc)$ use Retrieval algorithm from level-3 to traverse ID space with the return value list. If the returning list is NULL, forward to level-2. Else, For every element $d(\beta, loc)$ of the list, if it is not loaded, the query is processed by $d(\beta, loc)$; else, locate the content peer $c(\beta, loc)$ which has the content copy and redirect the query(o) to $c(\beta, loc)$ if it is online. $c(\beta, loc)$ will feedback the content object o to the client. If $c(\beta, loc)$ is offline, try the next one in the returning list up to the end of the list and forward the query(o) to source website.

Performance and analysis: Present experiments are conducted in a static environment where peers do not

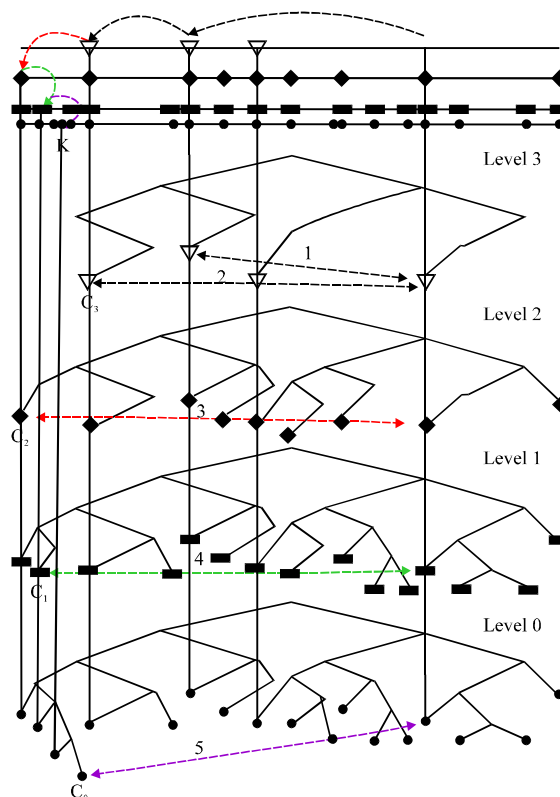


Fig. 3: Hierarchical directory peer DSHT

leave the system after joining it. Here, we focus on quantifying Hit ratio and Lookup latency in Aroma-CDN due to locality-awareness and interest-awareness. Hit ratio is the fraction of queries satisfied from the P2P system, an indicator of the degree of server load relief achieved, given that the fraction of queries reflected by the hit ratio are not redirected to the server. Lookup latency is the average latency taken to resolve a query and reach the destination that will provide the requested object(original server or content peer). Lookup latency is an indicator of the system's search efficiency.

We conduct simulation experiments using PeerSim (Jelasity *et al.*, 2010) that enables us to model the latency of each individual link. We generate an underlying topology of peers connected with links of variable latencies between 10 and 1000 msec. Localities are modeled using the binning technique. We use 4 layers DSHT and choose Chord overlay (Rao *et al.*, 2007) for its simplicity; we simulate its routing and adapt its key management service, to simulate the DSHT protocol. We assume that Aroma-CDN supports $|O| = 1000$ content objects which results in $|O|/8 = 125$ Dps. The population P is 1250 and underlying network nodes is 5000. We compare Aroma-CDN with




Fig. 4: Aroma-CDN’s lookup latency compared with Flower-CDN

Flower-CDN and DHT-Directory approach, Flower-CDN supports locality and interest but does not support DSHT, while DHT-Directory without any locality or interest considerations is widely employed in the P2PCDN literature (Stavrou *et al.*, 2002; Wang *et al.*, 2002; Ryu and Yang, 2005). We chose the Flower-CDN because it shares some similarities with Aroma-CDN with respect to the structured combined with unstructured architecture. This allows us to see the effects of locality-based petals and their gossip based management with DSHT support.

Present experiment measures the lookup latency. Fig. 4 shows the variation of the average lookup latency of a query with time: the lookup latency starts by decreasing and stabilizes around 150 msec shortly after the system warms up (i.e., less than 13 h), while Flower-CDN near 200 msec. Figure 5 shows the latency distribution of queries for both solutions: 92% of our queries are resolved within 150 msec while 87% of Flower-CDN’s queries take more than 200 msec. In Aroma-CDN, only first queries of new participants have to go through DSHT and result in long lookup latencies. Afterwards, queries are resolved within the local petal, achieving very short delays. In contrast, because of Flower-CDN’s slow D-ring query, it has longer lookup latency. Thus, we conclude that the locality-aware DSHT hybrid overlay of Aroma-CDN performs very well in providing efficient lookup.

In general, a smaller hit ratio means less queries are served from the P2P and instead go to the server. The Fig. 6 shows that the hit ratio eventually converges to 1 for the three solutions, convergence takes longer for Aroma-CDN compared with DHT but shorter than Flower-CDN because of content object-oriented key management, given that the search space is partitioned into petals. In fact, after 13 h, the hit ratio of Aroma-CDN is similar with DHT-Directory and after 24 h higher than that of

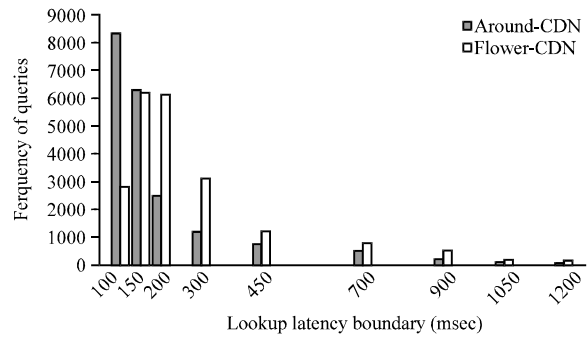


Fig. 5: Lookup latency distribution




Fig. 6: Hit ratio

Flower-CDN by 10%. This shows our solution can also have a good gain in lookup latency without hit ratio loss.

CONCLUSION

The study presents a new hybrid architecture based on P2P locality-awareness and interest-focusing in order to accelerate the client access the website content object and the content peer keep the content they retrieve and later to serve other peers that close them in locality. We combine DSHT with petals in order to get a faster query. Simulation-based evaluation shows that Aroma-CDN yields a high performance gains under large population.

REFERENCES

El Dick, M., E. Pacitti, R. Akbarinia and B. Kemme, 2011. Building a peer-to-peer content distribution network with high performance, scalability and robustness. Inform. Syst., 36: 222-247.

Freedman, M.J., E. Freudenthal and D. Mazieres, 2004. Democratizing content publication with Coral. Proceedings of the 1st Conference on Symposium on Networked Systems Design and Implementation-Volume 1, March 29-31, 2004, San Francisco, California.

- Iyer, S., A.I.T. Rowstron and P. Druschel, 2002. Squirrel: A decentralized Peer-to-peer web cache. Proceedings of the 21st ACM Symposium on Principles of Distributed Computing, July 21-24, 2002, Monterey, California, USA., pp: 213-222.
- Jelasity, M., A. Montresor, G.P. Jesi and S. Voulgaris, 2010. The Peersim simulator. <http://peersim.sf.net>.
- Rao, W., L. Chen, A.W.C. Fu and Y. Bu, 2007. Optimal proactive caching in peer-to-peer network: Analysis and application. Proceedings of the 6th ACM International Conference on Information and Knowledge Management, November 6-10, 2007, Lisbon, Portugal, pp: 663-672.
- Ryu, Y.S. and S.B. Yang, 2005. An effective peer-to-peer web caching system under dynamic participation of peers. IEICE Trans. Commun., E88-B: 1476-1483.
- Stading, T., P. Maniatis, M. Baker, 2002. Peer-to-peer caching schemes to address flash crowds. Proceedings of the First International Workshop on Peer-to-Peer Systems, March 2002, Cambridge, MA., pp: 203-213.
- Stavrou, A., D. Rubenstein and S. Sahu, 2002. A light weight, robust P2P system to handle flash crowds. Proceedings of the 10th IEEE International Conference on Network Protocols, November 12-15, 2002, Paris, France, pp: 226.
- Wang, X.Y., W.S. Ng, B.C. Ooi, K.L. Tan and A.Y. Zhou, 2002. BuddyWeb: A P2P-based collaborative web caching system. Web Eng. Peer-to-Peer Comput., 2376: 247-251.