Awake: decentralized and availability aware replication for P2P cloud storage

Yahya Hassanzadeh-Nazarabadi, Alptekin Küpçü and Öznur Özkasap
Department of Computer Engineering, Koç University, İstanbul, Turkey
{yhassanzadeh13,akupcu,oozkasap}@ku.edu.tr

ABSTRACT

The data availability of a peer-to-peer (P2P) cloud storage is negatively affected by the churn, which corresponds to the dynamic arrivals/departures of the nodes to/from the system. In this study, we propose a dynamic, fully decentralized, and availability aware replication algorithm named Awake with the goal of maximizing the availability of replicas in the structured P2P cloud storage. Our extensive analysis results show that Awake improves the availability of replicas with an average gain of about 21% compared to the best existing solutions, and it is independent of the churn behavior of the underlying system.

1 PROBLEM DEFINITION

A P2P cloud storage consists of a set of nodes where a node can be both a data owner and a data requester simultaneously. The data owner holds a set of data objects and wants to share them with a group of authorized nodes called its set of data requesters. To improve the availability of data objects under churn, the data owner makes copies of its data objects on the other nodes, which are called its replicas, and the procedure of replica placement is called replication. Once the replication is done, the data requesters can utilize the replicas. In the presence of replicas, failure of the data owner does not affect the availability of data objects. The traditional decentralized availability-based replication algorithms employ some kind of randomness in replicas placement. This prevents the algorithms to purely consider the availability patterns of nodes and causes poor availability of replicas. Likewise, all those algorithms make explicit assumptions about the underlying churn behavior of the system in replicas placement, which results in poor data availability in situations where the churn behavior of the system is mispredicted or changes over the time.

To improve the data availability of structured P2P cloud storages under churn, we propose a dynamic, fully decentralized, and availability aware replication algorithm, named Awake, that is independent of the churn behavior of the underlying system. In contrast to the decentralized counterparts, Awake does not consider any assumption about the underlying churn model. By employing Awake, a data owner can replicate its data objects in a fully decentralized manner such that the maximum availability of the replicas is achieved. The maximum availability of replicas is defined as the maximum average number of available replicas at a predefined time unit (e.g., an hour) of a fixed periodic time interval (e.g., a day).

2 AVAILABILITY AWARE REPLICATION

Awake is employed in a structured P2P system where the availability information of nodes is represented by an availability vector. Each cell of the availability vector corresponds to the availability probability of a node during a certain time slot of a fixed time interval. In this study, we consider the fixed time interval as a day and the time slots as hours. Nodes piggyback their availability vectors on the query messages they initiate or route. On receiving a message for routing, the node updates its availability view of other nodes using the received availability vectors. This piggybacking approach of availability information is common among all the best existing decentralized solutions [2]. When a data owner wants to place its replicas, it executes Awake on the set of availability vectors that it has obtained. Awake models the availability aware replication as an integer linear programming problem, solves it on the input set and outputs the replica set to the data owner. The data owner then shares the obtained replica set with its set of data requesters. Our recent work presents the Awake model [2]. In the case that the underlying churn behavior changes and the availability of replicas drops significantly, the data owner re-executes Awake using its own updated availability view of the system.

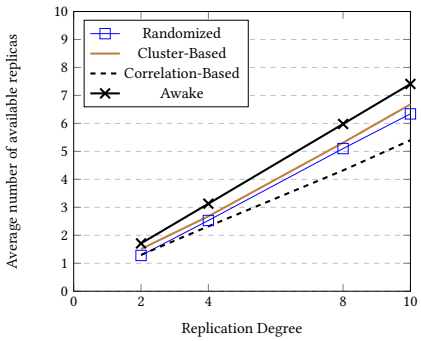

Figure 1: Average number of available replicas at each hour vs replication degree.

We implemented Awake as well as the state-of-the-art availability based replication algorithms in the SkipSim [1] simulator. We simulated each algorithm on 100 randomly generated topologies each with 128 nodes examined under high, moderate and low available churn models. Figure 1 shows that Awake improves the average availability of replicas with the gain of about 24% compared to the best existing solution under the high available churn model. Likewise, Awake outperforms the existing solutions under the moderate and low available churn models with the gain of about 14% and 26%, respectively. The simulation results indicate that for a certain number of replicas and a fixed churn model, Awake performs independently of the system size. In the poster presentation, the design and algorithmic details of Awake and the extensive experimental results would be provided.

REFERENCES

[1] SkipSim: git@gitlab.com:yhassanzadeh13/skipsim-distribution-bundle.git.
[2] Yahya Hassanzadeh-Nazarabadi, Alptekin Küpçü, and Öznur Özkasap. 2016. Awake: decentralized and availability aware replication for P2P cloud storage. In *Smart Cloud*. IEEE.