

Churn in *psiX*

Lokesh Bindal, Steven Lewis, and Vasil Slavov
{lbvd7,sdlf2b,vgslavov} @umkc.edu

1 Setup

Network churn is simulated by continually joining and disconnecting peers from the DHT. To indicate how often nodes will join or leave *psiX* per minute (that is, how many churns per minute are desired), an integer i is given. Churns happen on some random interval ranging from 1 second to $2i$ seconds. This allows the average rate of churns to reach a mean rate of i churns per minute. For example, for 5, 7, 9, and 12 churns per minute, the average time between churns is 12, 8.5, 6.5, and 5 respectively. After a random interval, a peer is selected at random from the list of all available peers. If that peer is already a part of the DHT, the peer is stopped; otherwise, it is started and added to the DHT. The list of active peers is continually updated to mirror the current state of peers in the DHT. Initially, *psiX* was started on 55 PlanetLab nodes. Due to churn, the number of active hosts varied between 20 and 55. The average number of peers was 30.

Publishing and querying was done on the bootstrap node only. A total of 3459 XML documents were published (about 100 documents for each of the 35 DTDs in /home/raopr/cs5590ld/P2PXML/DTDs). A total of 703 XPath queries were submitted (about 20 per DTD).

The default replication rate in *psiX* is 12. In order to determine the effect of the replication rate on churn, publishing and querying in *psiX* was tested with three different replication rates: 6, 12, and 18. The effect of the lookup timeout on churn is left for future work.

2 Observations

2.1 Publishing

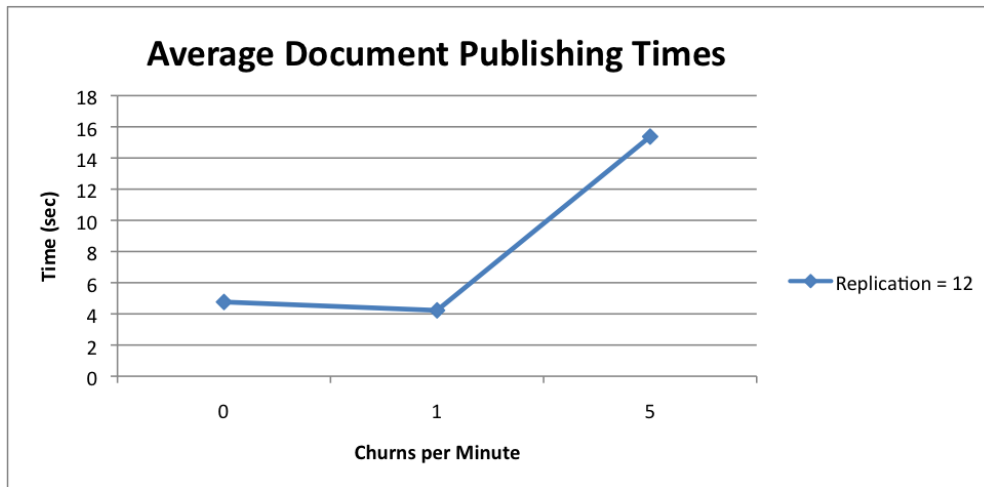


Figure 1

As is evident from Figure 1, high churn causes publishing times to increase rapidly. At the default replication of 12, the average publishing times grew from around 4.2 seconds to 15.3 seconds. Due to time limitations and the evident effect shown above, no higher churn testing was performed.

Figure 2 shows the different publishing times for the three replications under 1 churn per minute. The high publishing time for replication 6 may be a result of the loss of replicas due to churn.

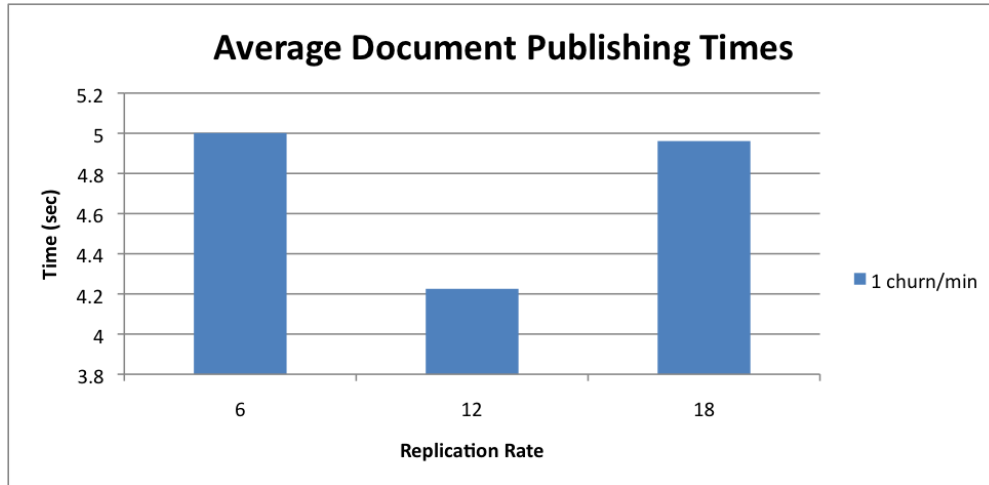


Figure 2

2.2 Querying

The amount of data lost as a result of network churn is directly correlated to the query response times. If a primary replica of a key-value pair is lost due to a peer leaving the DHT network, then the next predecessor peer that contains a replica must be contacted. If that peer has also left the DHT network this process will continue until either a peer that hosts the data we are querying for is found or the lookup operation times out. Consequently, a lower number of replicas results in a higher probability that the data is completely lost and the lookup operation will time out. The higher replication rates cause the DHT to attempt to contact more peers in an effort to locate the correct data.

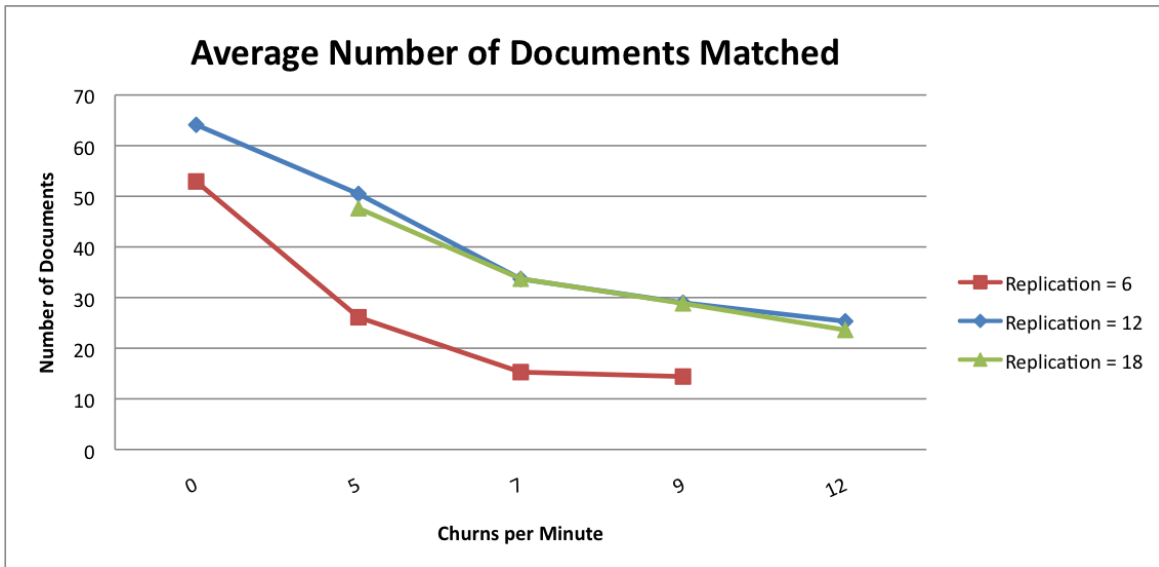


Figure 3

As seen in Figure 4, in the case of 6 replications, the query response times were actually higher than the default replication rate. In some instances, since there are fewer replicas, the effect of churn may have completely eliminated all peers which store those replicas, resulting in an increased response time due to the query failing after a timeout period. In addition, Figure 3 shows that the average number of documents returned is lower due to the complete loss of replicas.

Churn does not seem to affect the number of matched documents when the replication rate is high (Figure 3). However, at high replication rates, query response times are higher than the response times for the default replication rate (Figure 4).

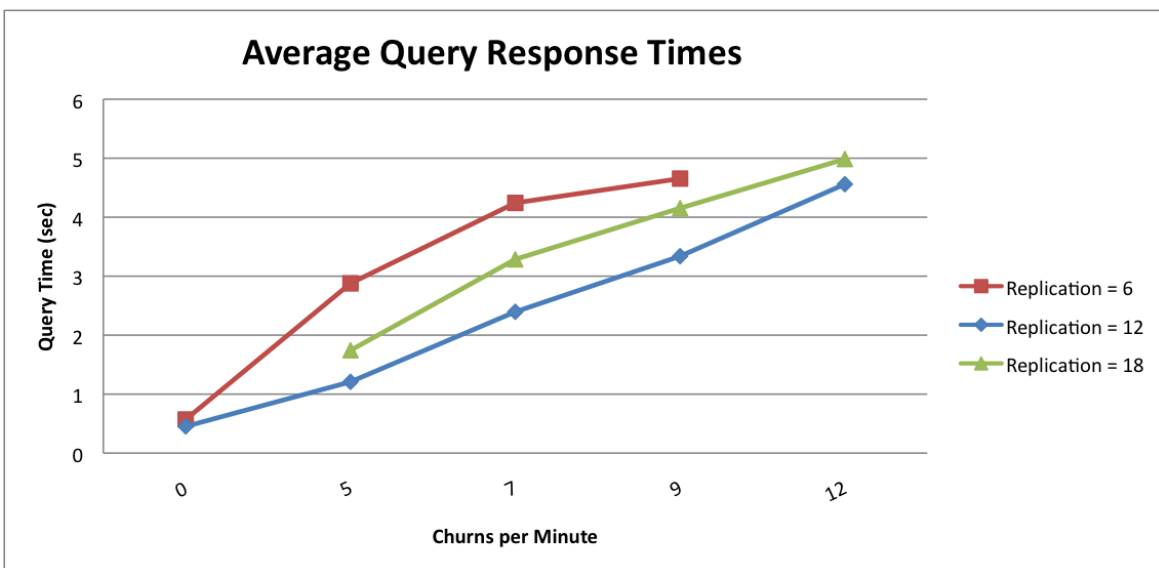


Figure 4

Considering the above results, a replication rate that is too high or too low can result in higher query response times with no added benefit of a higher percentage of documents matched. The default replication rate of 12 produced a comparable number of average documents returned and did so with a lower average query response time. Of the tests performed, the replication rate of 12 had the best performance. The values chosen for testing represent an extreme variation on the replication rates and can be further fine-tuned to find the most optimal setting. Note that the optimal replication rate will vary as the scale of the DHT changes.