Urn models and peer-to-peer file sharing

Ilkka Norros and Hannu Reittu
VTT Technical Research Center of Finland
e-mail: firstname.lastname@vtt.fi

We consider an open network, with constant rate of incoming 'peers'. A peer can contact and communicate with any other peer in the system, corresponding to an Internet's overlay, where knowledge of peer's IP-address enables communication with it. In this system one peer, the 'seed', stays in the system and holds a file and wishes to distribute it to all its peers. One way of doing this, and suitable for large number of peers, is the idea of Bit-Torrent like systems: as soon as a peer receives the file, it becomes a seed itself. To enhance performance, the file is divided into small chunks that are spreaded in a similar fashion.

In our previous work we considered a system without incoming peers, see the cited paper [1]. The system started in a state with seed and large number of peers without chunks. Then making random contacts the chunks start to propagate in the system. It was noticed that in the case of 'selfish' peers that leave the system immediately after receiving the entire file, the copying process is asymptotically, with increasing number of peers, equivalent to Pólya's urn model, with characteristic random proportion of each chunk in the system. This imbalance leads to the 'rare chunk phenomena': one chunk is not able to become common and as a result forms a bottleneck of performance. In an open variant of this problem, with incoming peers, this could lead to instability, the number of nodes in the system would grow unboudently, since more and more peers are searching for the one missing chunk, the rare chunk.

A relevant question is whether it is possible to avoid such imbalance and instability without using centralized enforcing? One source of ideas could be models related to physics, see the cited paper [2]. Indeed, so called Ehrenfests' urn model gives almost an ideal balance in a closed system. Another example is so called Friedman urn, with analogous result for an open system with flow of incoming particles. The first approach, already analysed in the cited paper [1] , is as follows. The system starts with two nodes with chunks 0 and 1 each. Rest of nodes, very large in numbers, have no chunks and make random contacts, independently of each other and they obtain a chunk without delay. The rule is that if an empty node first contacts the node having chunk 0 (1) it decides to get the opposite chunk 1 (0) first. Then number of chunks is very likely balanced all the time. Indeed, it was proved that this fraction converges almost surely to $\frac{1}{2}$ as the number of nodes goes to infinity. In open system upon arriving in the analogous system, the arriving peer makes a uniformly random contact with nodes having at least one chunk. If the corresponding node has, say, chunk 0, it does not copy it, but decides to find chunk 1 first. After this entrance phase it starts to search for the chunk 1, and then the chunk 0, upon finding that the node leaves the system.

We furnished this type of models and analysed their stability using computer experiments and analytical approach. We found that indeed, such approach can give a stable system whilst without such mechanism, the number of peers in the system grows in the pace of arriving nodes. However, the dynamics of such random system seem to be untrivial, showing strong, and more or less regular oscillations.

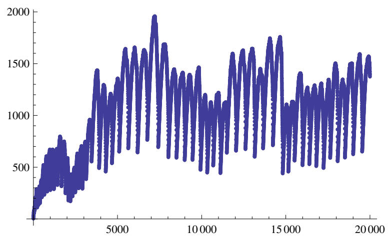

A sample path of number of nodes in the system, after starting from seed alone, arrival rate is 10 new peers / download time.

References:

[1]I. Norros, B. Prabhu, and H. Reittu, Inter-Perf, 2006, Pisa.

[2]R. Pemantle, Probability Surveys, Vol 4 (2007), pp. 1-79.

WIOPT 2008, 1st – 3rd Apr 2008, Berlin, Germany.
 ISBN 978-963-9799-18-9
DOI 10.4108/ICST.WIOPT2008.3172