

Deploying VoDKA in a Network of Peer-to-Peer Public Information Screens ^{*}

Carlos Abalde, Víctor M. Gulías, and Laura M. Castro

MADS Group. Computer Science Department
University of A Coruña (Spain)
{cabalde, gurias, lcastro}@udc.es
<http://www.madsgroup.org>

EXTENDED ABSTRACT

Research on content distribution networks (CDNs) pushes hard to find novel approaches in publishing a huge quantity of relevant information to be accessed by a large number of clients. This paper presents our results building a peer-to-peer (P2P) video CDN developed as an extension of an existing on-demand streaming server (VoDKA [1], Video on Demand Kernel Architecture). It has been designed to improve streaming server capacity without any infrastructure upgrade when it must deal with big networks of public information screens accessing it. In the final paper, additional background information about the problem and VoDKA is given.

Intuitively, a P2P approach is better suited to deal with mass-scale content distribution when common access patterns to the contents can be identified among neighbor —locally connected— nodes. This is the case of a network of public information screens. Obviously, a P2P approach also involves more effort in terms of coordination, resource management, heterogeneity... than a traditional client/server approach [2]. In the final paper, some comments on up-to-date P2P content distribution technologies are given.

Minimize coupling, design a scalable and fault tolerant architecture, allow efficient screen departure and arrivals, integration with an existing client/server CDN and minimize deployment and maintenance costs are the main design goals of our video CDN. Based on the survey of P2P content distribution technologies found at [3], we conclude that a fully decentralized structured architecture is the best design choice. Therefore, we devise a fully decentralized structured architecture based on a distributed hash table algorithm (Chord [4]) and integrated with our streaming server.

The streaming server becomes a central storage in charge of screen priorities and resources management. Screen requests will be multiplexed on TCP/IP connections and scheduled in order to reach program grid deadlines.

Screen behavior is built on top of a distributed hash table which is used as a state dissemination infrastructure. Then a layered architecture is defined and a P2P screens coordination algorithm designed. Both topics are extended in the final paper. On the other hand, the P2P screens networks and the streaming server

^{*} Partially supported by MEC Project TIN2005-08986.

are easily composed through the VoDKA traders architecture. A screen/location priority framework is defined as the basis for our resource scheduler algorithm. This algorithm is responsible for optimal resource usage and deadline achievement, and it is also extended in the final paper.

As shown by the benchmarking results (figure 1), our strategy solves the proposed problem without any extension of the central architecture. A set of close-by screens —information screens on the same LAN— are managed by the VoDKA system as about 1.2 screens independently of the set size.




Fig. 1. LAN-screen equivalence

Therefore, as shown in our experiments, the design is an appropriate approach for highly-correlated mass-scale content distribution, with interesting scalability and availability features. In the final paper, some additional benchmarks are given.

References

1. Gulías, V.M., Barreiro, M., Freire, J.L.: VoDKA: Developing a video-on-demand server using distributed functional programming. *Journal of Functional Programming* **15**(4) (2005)
2. Padmanabhan, V., Wang, H., Chou, P., Sripanidkulchai, K.: Distributing streaming media content using cooperative networking (2002)
3. Androutsellis-Theotokis, S., Spinellis, D.: A survey of peer-to-peer content distribution technologies. *ACM Computing Surveys* **36**(4) (December 2004) 335–371
4. Stoica, I., Morris, R., Karger, D., Kaashoek, M.F., Balakrishnan, H.: Chord: A scalable peer-to-peer lookup service for internet applications. Technical Report TR-819, MIT (March 2001)