Tech. Topic 1: Content Distribution Networks

I. Introduction

In this week's technical topic on Content Distribution Networks (CDN's), we have two papers[1], [2] and a patent[3] as part of the reading list. In this report, we choose to present a review of the patent in detail [3]. This paper is chosen, because it is the beginning of a rising content distribution industry. It is interesting to review the beginnings of this technological contribution.

II. Paper 1 Review: Global Hosting System

The authors of this paper, F.T. Leighton and Daniel M. Lewi are founders of Akamai, a content distribution network industry. This patent is held by Akamai technologies.

Currently, Akamai distributes over 20% of the Internet traffic over its platform. This patent marks the beginning of Akamai as a company with patented technology and services. At the time the patent was published, looking forward, they were able to achieve the goals that they had initially set out to achieve, and now provide useful services on the Internet.

This patent is very clear, comprehensive, and self-explanatory. It is structured in a way that is easy to follow and to maintain an understanding of the thinking process involved in its proposition (problem, followed by purpose of current invention, then a detailed description of operation); the authors move from one point to the other in a logical manner. The patent, however, is lacking real-life data or statistics that could have been useful to discuss the, at the time, current state-of-affairs and to motivate the need for dedicated content servers managed by a particular organization (which in this case is Akamai [1]). However, it can also be argued that the need for the proposed technology is clear, and the supportive statistics or data should be intuitive.

They begin by discussing how content providers currently (at the time) distribute content over the Internet, and why existing solutions are neither economically feasible nor scalable, and how they require high maintenance costs from organizations that are dedicated towards other business paths and might not be willing to, or at least have to, incur such high costs. The authors describe how existing content distribution methods have negative effects on advertising (content providers are unaware of user demand for their content, ex. by not being able to maintain a hit count on their webpages), user-satisfaction (distributing outdated information to users), and, as a result, reduced revenue. Therefore, they begin by providing an excellent motivation for their decentralized hosting solution while keeping in mind the need for the content provider to maintain control over the delivered content. The proposed solution clearly presents benefits to content providers and Internet users.

The authors follow with a description of the objectives or purpose of the invention, which are practical, useful, and reasonable. They appeal to the needs of content providers as well as the end-users. This builds-up the argument further for their invention and establishes the need for it in addition to the motivation for it.

It is important that they keep in mind the content provider's need to retain control of the content. They also address memory and replication concerns, and how to maintain efficiency.

When assigning objects to servers, however, it seems that they do not take into consideration the load that might be imposed by that object request; however, it later shows that the invention addresses this issue, which should have been briefly mentioned earlier.

The presence of a hierarchy, and particularly a two-level hierarchy, is a good idea to maintain low overhead, low cost, as well as load balancing.

The authors later move on to a detailed description of the invention, which, as mentioned earlier, is well-written. The figures referred to were clear, useful, and relevant to understanding certain concepts, and the explanation of the figures was easy to follow and understand. Number-labeling the components of the diagrams made it easier to follow with the explanation.

Their solution of allowing the content provider to send its main HTML page with modified URL's, where the modified URL's are for those objects that will be fetched from the CDN's, is intelligent and practical. To further encode the modified URL's so as to include information about the object, as well as to use embedded strings so as to attain load-balancing and proximity of source content (by using the top and low level DNS servers), is similarly an intelligent and practical approach to achieving the objectives of this invention.

They discuss how it could be possible to have multiple levels in the DNS hierarchy in the porposed invention, however, they mention it very lightly. Essentially, in this patent, there are two DNS hierarchical levels that are of concern. Any more levels will add unnecessary complexity.

CDN's host their infrastructure in Internet Service Providers (ISP's). CDN's are, therefore, practical and implementable because they utilize already existing infrastructure; the deployment of CDN's starts on well-maintained and reliable ground. They later discuss how the ISP's will determine where the hosting servers will be placed in their network, and this will be

[1] http://www.akamai.com/

done through monitoring network traffic, or requests. It is unclear, however what the ISP's incentive would be to do this. Will the ISP be financially compensated by the companies hosting the CDN's?

The step-by-step, chronological, description of how the request by the client is processed is extremely useful and creates structure in the reader's mind.

When storing content, one concern would be staleness of data. However, they porpose a method of updating the stored content which only re-fetches the modified objects from the content provider so as to save bandwidth.

When explaining the use of the two hashing functions, the explanation got a little confusing.

This invention requires minimal synchronization between the content providers and the host servers, and this is performed in the initial stages when content providers are accepting the service provided by CDN's. This synchronization can be thought of as an initial contract with the CDN. In reality, now-a-days, a contract is maintained by Akamai with its customers where the contract is configured towards the needs of the customer.

Their method of determining host location by using the IP of the client, or that of the ISP DNS server may not be a completely accurate measure. A client can be located far from the ISP DNS server, or it could be behind a NAT. They later further address the IP-locating issue and propose an additional safeguard, which is by examining the IP address of the client. However, this safeguard is unreliable and unnecessary; another more effective solution is required, if any.

Furthermore, when assigning servers to a particular IP block (location-based server-pool selection), this assignment is maintained/updated continuously, yet: how often is continually, and what triggers a change in this mapping? This, however, is a good approach because, if the top level DNS servers only pick the low-level servers which are close to the client, then this might be a concern in case of regional outages (i.e. all ghost servers in a particular region could be overloaded or down). Since the assignment of low-level DNS servers to clients is more flexible, the invention is able to overcome the problem of regional outages.

It is clear that delay is the main concern, i.e. minimizing delay, however, not at the cost of load. Their method for load-balancing, as performed by the low-level DNS servers by monitoring the various ghost servers, is simple and effective. Load-balancing is performed continuously and is maintained.

They address redirecting the client to different servers in case of increased load. This approach involves running a software in the client's browser (or media player) that monitors the performance of the client's connection and, potentially, that of the status of the network. This solution requires cooperation by the clients and processing overhead; a lot has to be taken into consideration, such as: what is the redirection period (how often do we redirect)? What is the performance threshold after which we should switch? Does it depend on the size of the object we are transferring? How much delay is incurred and is it worth it? What is the tradeoff.

The interesting part of this invention is that there are a lot of parameters involved in building this system which can only be given value from real-word experimental analysis. Actual testing and deployment is necessary to assign values to these parameters (such as, what is the threshold for object requests? How will cost matters be determined?).

The system is built with high fault tolerance. If a ghost fails, it has a buddy ghost to rely on. If ghost servers in a location go down from regional load, top level DNS servers will assign a different ghost server location to clients.

Breaking down and discussing the advantages afforded by the suggested global scheme are accurate, clear, and descriptive.

Having worked at Akamai, the objectives and advantages discussed in this patent are being experienced perfectly. Akamai provides services to a wide range of content providers, and in the past years, has hosted and facilitated major flash crowd events (such as the recent presidential elections). One might argue that statistics of the current (at the time) traffic would be useful, such as what percentage of traffic is static versus dynamic. However, the need for a content distribution network and its benefits are apparent; there is no need for live traffic data to support it.

There are a few spelling mistakes in the patent, missing letters.

REFERENCES

[1] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy, "An analysis of internet content delivery systems," *SIGOPS Oper. Syst. Rev.*, vol. 36, no. SI, pp. 315–327, 2002.

[2] Z. Al-Qudah, S. Lee, M. Rabinovich, O. Spatscheck, and J. Van der Merwe, "Anycast-aware transport for content delivery networks," in *WWW '09: Proceedings of the 18th international conference on World wide web*. New York, NY, USA: ACM, 2009, pp. 301–310.

[3] F. T. Leighton and D. M. Lewin, "Global hosting system." U.S. Patent No, 6,108,703, August 2000.