

CP-CDN: Content Delivery Powered by Context and User Intelligence

Abhishek Sharma

Asst. Professor, Department of Comp. Sc. & Info. Tech., Graphic Era Hill University,
Dehradun, Uttarakhand India 248002

Article Info

Page Number: 155-163

Publication Issue:

Vol. 70 No. 1 (2021)

Abstract

In order to improve end users' Web access experiences, today's Internet content providers generally employ Content Delivery Networks (CDNs) to provide material to end users. To comprehend reasons, methods, and amounts of CP-level intelligences may aid content distribution, we employ a measurement-based approach. The largest content provider in China, Tencent, constructed a CDN, which we first offer as a measurement study. Contrarily, CPCDN may leverage context intelligence—which captures how contents are constantly manipulated then synthesized in various circumstances—in order to optimise the user experience and service quality throughout content delivery since it is aware of the precise environment in which material is consumed. In this study, we use a measurement-based approach to reasons, methods, and amounts of CP-level information may aid content delivery in order to support the quick transition to CPCDN. We proactively replicate content based on interest groups based on crowd patterns so that users can access their desired content from nearby CPCDN servers. CPCDN is able to raise the level of both content delivery efficiency and user experience by using the potent user intelligence that captures how contents are constantly exchanged and disseminated across social networked users.

Keywords: Content Delivery Networks, Measurement-Based Approach, Social Network, Context and User Intelligence.

Article History

Article Received: 25 January 2021

Revised: 24 February 2021

Accepted: 15 March 2021

1. Introduction

From people who unintentionally utilise the cloud to get their favourite songs to businesses that purposefully make advantage of its potent resources, the cloud has already altered business and daily life. A perfect storm of opportunity has been created for businesses to use the power of the cloud to optimise, innovate, and disrupt business models thanks to recent technological and social networking developments. According to a poll of 572 business and technology leaders worldwide, while cloud computing is universally acknowledged as a crucial technology, only a small number of organisations are actively utilising it to foster the development of new business models. Performance and scalability are key problems for websites as Internet traffic increases and websites get more sophisticated. In order to give website users dynamic, interactive, and personalised experiences, dynamic content generating apps are being used more often. However, each request necessitates computation and communication between several components.

According to recent research and analysis, there are around 800 million Web pages that can be publicly indexed. This essay provides statistical analysis of the frequency and kind of Web page alterations as well as early data on such adjustments. It makes an estimate of how

frequently search engines must reindex the Web in order to stay up to date. Online social networks (OSNs) have additional design issues as a result of their scalability challenges, which sometimes necessitate expensive re-architecting. SPAR is a middleware for social partitioning and replication that transparently uses the structure of the social network to achieve data localization while reducing replication.

This strategy has several benefits, including the opportunity for application developers to assume local semantics, scalability with the addition of commodity servers with little memory as well as network I/O needs, along with redundancy at a significantly reduced cost. Without altering a single line of its application logic, SPAR has minimal overhead and may assist a well-known open-source Twitter clone in scaling to Twitter's level.

In addition to location awareness, context awareness is a feature of mobile devices. It has been used in business theory in connection to contextual application design and business process management concerns. Its roots are in ubiquitous computing. We suggest a context-aware system for information integration to close this gap. It can handle various information kinds and assess their applicability to the user's present situation. In order to create a more universal and adaptable system that automatically selects the most relevant material and provides it to the user in a compressed format that facilitates ease of consumption and understanding, we log user interaction and perform usage mining using OLAP to uncover context-dependent preferences for various information types.

Although the advent of ubiquitous computers and communications has long been predicted, there is still a big difference between what consumers are experiencing now and what is projected. That is not to downplay the fact that there has been a lot of development. As a more in-depth grasp of the practicalities develops, it instead serves as evidence of the complexity of system design and execution. Thus, even in research labs, fully ubiquitous computer systems are uncommon. Therefore, the general public and to a lesser extent researchers can only make conjectures about what a world dominated by computing technology will like and how it would operate.

Users can get a preview of what may be in store in the future thanks to mobile computing, one of the offshoots of the ubiquitous computing paradigm. It's true that the technology and uses are constrained. Even though mobile computing will always be inferior to fixed networked workstations, it is possible that over time this scenario will improve. This is especially true, of course, in the case of multimedia applications. The implementation of robust multimedia-enabled apps and services has been negatively impacted by the constrained processing resources and the relatively weak bandwidth of cellular data networks.

In the framework of ubiquitous computing, this chapter examines a number of uses of semantic technologies include context management, adaptive human-computer interaction, privacy protection, including social communications. It suggests many lines of inquiry for the development of a unified omnipresent system based on semantic and social web technologies. The growing complexity of the information society hurts users. at a time when the amount of information available online is increasing dramatically. It is necessary to expand the traditional understanding of human-computer interaction to take into account a cloud of interconnected,

heterogeneous electronic devices that provide services to the user. Mobile phones and PDAs are examples of small mobile devices with wireless capabilities that are excellent for providing users with access to ubiquitous communication services.

The bulk of office workers, however, continue to spend the whole day glued to their desktop email applications. A means of filtering the messages vying for the user's attention is required to deal with this. Context-aware systems may infer a variety of information from their surroundings, such as who is there or what job is being carried out, that can guide their action. In these systems, decisions can only be taken if there is a sufficient amount and quality of data to support them. A network of sensory devices must be present, each of which offers unique inputs that add to our understanding of the system as a whole.

2. Literature Survey

A simple and totally automated method for locating and categorising content storage and delivery infrastructures is called web content cartography. In accordance to DNS measurements plus BGP routing table snapshots, it leverages these metrics to find hosting infrastructures. The main aim is to gather the IP addresses that DNS delivers for different popular and uncommon hostnames from various geographic perspectives. To identify between various hosting infrastructure deployment methodologies, features like prefix, AS, and location of an IP address are employed. We distinguish between various kinds of hosting infrastructures using both the granularity of BGP prefixes and /24 subnetworks. There are counter-examples, such as Meebo and Netflix, but our strategy takes them into account by clustering hostnames differently. Our clustering method succeeds in locating hosting infrastructures in the wild and can distinguish amongst hosting infrastructures if they are looked after by different administrative entities. Our measuring method depends on volunteers to run a programme that asks the locally configured DNS resolver, a Google Public DNS resolver, as well as an Open DNS resolver for more than 7400 hostnames and saves the whole DNS answers in a trace file [1].

This study examines the characteristics and relative effect of 1.6 million Twitter users in 2009 by monitoring 74 million diffusion events over a two-month period. It discovers that people with a high number of followers and a history of influence tend to produce the greatest cascades. Additionally, URLs were more likely to spread if Mechanical Turk gave them higher curiosity ratings and/or when they generated favourable reactions. Predictions of which specific individual or URL will produce significant cascades, however, are not always accurate. The most crucial information in this article is that influencer ex-ante prediction is crucial for word-of-mouth marketing as well as comparable operations. We followed the dissemination of the URL from its origin at a certain "seed" node via a series of reposts until the diffusion event, or cascade, stopped in order to determine the impact score for a given URL post. We have three options for how to allocate the relevant impact if B has many friends who have previously posted the same URL: full credit to the friend who posted it first, full credit to the person who posted it most recently, or divide credit evenly among all earlier-posted friends [2].

The combined issue of resource supply with caching (Specifically, replica placement) for cloud-based CDNs is addressed in this study with a focus on addressing dynamic demand

patterns. It comprises of a provisioning and caching algorithm framework known as Differential Provisioning and Caching, which leases cloud resources to create CDNs and caches contents in order to minimise the overall rental cost while meeting all needs. In order to maximise the overall demands met by the supplied resources, the Caching and Request Balancing (CRB) algorithm constantly modifies the content placements underlying route tables. Under rapidly fluctuating demand patterns, the suggested algorithms are effective in terms of both performance and leasing cost [3].

This article offers a definition as well as examine PORs, or proofs of retrievability. With the use of a POR scheme, an archive or backup service (prover) is able to provide a clear demonstration that a user (verifier) are capable of restore a target file F , i.e., that the archive preserves and consistently communicates adequate file data for the user to thoroughly reconstruct F . The usage of sentinels in a POR is comparable to other methods, such the "ringers" utilised, that predicated upon the encoding of secret check values in files. The POR techniques put out here stand out due to the way they combine cryptographic sentinels concealing with error-correction to increase the efficacy of spot-checking for the unique instance of file verification. Although Lillibridge et al.'s solution assumes different MACs on each block and fails to specifically tackle error-correction for the single-server situation, they do propose applying error-coding to a file together with spot-checks of file blocks carried out by system peers. NR (Naor and Rothblum) put out a more theoretical relevant finding. The authors provide a technique in which blocks are MACed after an error-correcting code is applied to file F . They also demonstrate security limitations on their architecture and offer a straightforward, formal definition of security. The suggested approach does not address the situation because it presumes the verifier has unrestricted accessibility to the encoded file F , whereas a server reports file blocks in F erratically.. Given the growing popularity of storage-as-a-service and users' increasing reliance on third parties to keep sensitive data, the authors feel that the sentinel-based POR protocol is suitable for real-world implementation. Construction of a secure storage system can also be aided by contractual and legal safeguards [4].

The technique developed by Sebe et al. is modified in this article to provide public verifiability without disclosing any personal information to independent verifiers. The suggested probabilistic checking approach, which only tests the integrity of c blocks, is used to obtain the probabilistic assurance of data integrity. The suggested protocol has been shown to be private and secure against outside verifiers and servers. We are still developing the protocol's ability to accommodate data level dynamics. Forward error correcting codes are used by Curtmola et al. to increase the PDP protocol's resilience. Wang et al. present a protocol for providing remote storage security in the setting of many servers, meanwhile Ateniese et al. give an environment that may generate a public verifiable proof of storage system from any public key homomorphic linear authentication mechanisms. For the multiple replica situation, Hao and Yu offer a remote data possession verification protocol that supports public verifiability and privacy from outside verifiers [5].

3. Proposed System

In the suggested system, the structure of a CPCDN, where content context and user information may be utilised for more intelligent content distribution, is presented, pushing content providers to construct their own CPCDNs. A vast dispersed network of computers set up in several data centres all over the Internet is known as a content delivery network (CDN) or content distribution network (CDN). High availability and high performance content delivery to end users is the aim of a CDN. A significant portion of the material on the Internet nowadays is served via CDNs, which involves web content (text, graphics, even scripts), content that can be downloaded (media files, software, and documents), applications (e-commerce, portals), live and on-demand streaming of content, and social media. In order to reach their end-user audience, content providers like media corporations and e-commerce merchants pay CDN operators. ISPs, carriers, and network operators get compensated by a CDN for housing its servers in their data centres. Along with improved performance and availability, CDNs also offload traffic provided directly from the origin infrastructure of the content provider, potentially saving the latter money.[1] Additionally, CDNs offer some defence against DoS assaults to the content provider by absorbing the attack flow through their extensive dispersed server infrastructure. There has been a recent tendency to employ a hybrid strategy that incorporates P2P technology, although the majority of early CDNs provided content via dedicated servers that were maintained by the CDN. Dedicated servers as well as other peer-user-owned machines may be used to provide content in the hybrid model, depending on the situation.

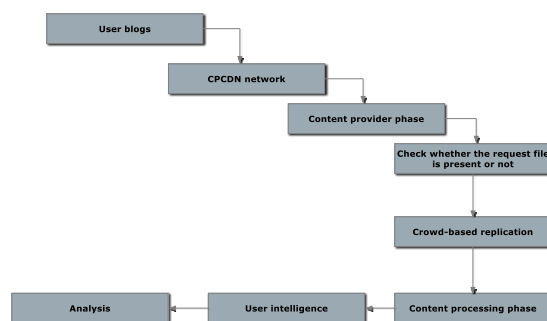


Fig 1: System Architecture

On the basis of context and user intelligences, we provide CPCDN delivery techniques. In order to better allocate network resources, we allowed the CPCDN to forecast how well-liked contents would perform among users of social networks. On the one hand, this helps us Understanding the importance of contents when they are presented in diverse situations can help you build the content replication with user request schedule techniques. The goal of this research is to demonstrate that using CP-level intelligence for content distribution can really have advantages. We will then discuss our measurement experiments that shed light on the possibility of a CPCDN.

A CDN improves speed by caching or copying material across certain mirrored Web servers (i.e., surrogate servers), which are carefully positioned at various places for the purpose to deal with the sudden surge in Web content requests, commonly known as the flash crowd or

SlashDot effect. The users are routed to the closest surrogate server. This strategy lessens the influence of the network on how quickly user requests are answered. When referring to digital data resources in the context of CDNs, the terms "content" and "metadata" are often used interchangeably. Static, dynamic, and continuous media data, like audio, video, documents, photos, and Web pages, are all included in the encoded media. The content description known as metadata enables the identification, discovery, and administration of multimedia data as well as makes it easier to analyse multimedia data. Content might be persistent or transitory data stored within the system, pre-recorded or obtained from live sources. The fundamental reference model for Open Systems Interconnection (OSI) may be thought as as receiving a new virtual overlay from CDNs. This layer offers overlay network services that rely on application layer transport protocols like HTTP or RTSP. The content supplier, the CDN provider, and the end consumers are the three main pillars of CDN architecture. One who delegated the URI name space of Web objects to be delivered is referred to be a content provider or customer.

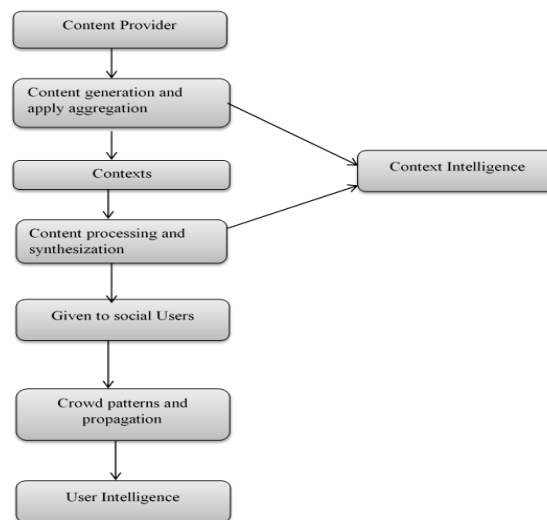


Fig 2: Flow Diagram

These items are stored on the content provider's origin server. A CDN provider is a private organisation or business that offers infrastructural services to content providers so that material may be delivered punctually and reliably. Entities that access material from a content provider's website are known as end users or customers. To replicate material, CDN providers utilise cache and/or replica servers positioned in many areas. Edge servers and surrogates are other names for CDN cache servers. We shall interchangeably use these words throughout this essay.

Due to the prompt creation of the components, CPCDN, which is aware of the delivery contexts, may enhance the content delivery by giving various components a higher priority for distribution. The CPCDN may lead users to obtain crucial materials from fast peering servers thanks to the component prioritisation. Context-aware content distribution is made feasible by a CPCDN because it has access to user and content information (such as content composition and user profile) and can track network information (such as the latency and bandwidth between peering servers and users). Users may concurrently download all of the webpage's components thanks to CPCDN, which also establishes the relative importance of the various content elements. The following benefits of the suggested strategy are listed:

- Our measuring study also demonstrates that specific social behaviours among users of online social networks have a significant influence on whether individuals would find particular contents popular, if at all.
- Context-aware and social content distribution on traditional CDNs may be greatly enhanced by the concept of CPCDN.
- The CPCDN is able to make use of user intelligence and context intelligence to enhance comprehension of upper-layer content features while also optimising under-layer content delivery tactics.
- The CPCDN technique may schedule the requests to servers effectively based on their capacity and traffic, preventing servers from being overloaded.

4. Results

A huge dispersed system of servers set up in several data centres all over the Internet is known as a content delivery network (CDN). This research suggests a measurement-based method to comprehend the reasons, the means, and the extent to which CP-level intelligences can facilitate content distribution. To unload traffic from the origin infrastructure of the content provider and offer end users with material that is highly available and performant, CDNs are deployed. P2P technology has been used in hybrid models in recent trends. By caching or copying material among mirrored Web servers (surrogate servers), CDNs improve speed.

Context-aware content distribution is made feasible by CDNs' ability to prioritise the delivery of various components of content and keep track of network data. As a means for SPs to increase income and cut expenses, CDN federations are gaining popularity. Measurement-based analyses show that the CPCDN design is successful in terms of decreased load latency, enhanced local download, and enhanced social media content prediction.

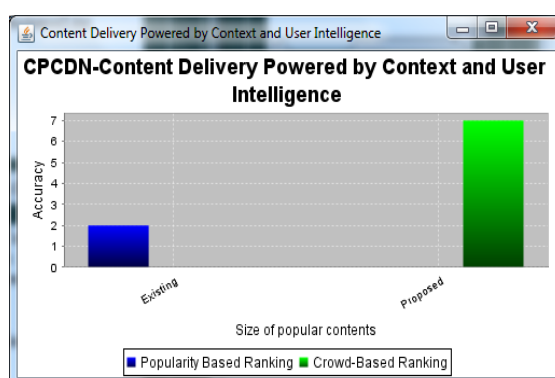


Fig 3: Performance Analysis

5. Conclusion

As the amount of traffic on networks continues to rise, CDN federations are becoming more popular as a means for SPs to make money and cut expenses. The core ideas and principles necessary to make CDN federations a reality have already been confirmed by Cisco's work with top SPs. As this study moves forward, CDN federations will soon be a practical method for Ps to obtain the proper commercial rewards for their essential role to the value chain for

content delivery. Our measurement-based evaluations show that our CPCDN design is more effective than traditional CDN approaches in the following ways: context-aware request scheduling shortens the time it takes for a webpage to load; crowd patterns are taken into account when replicating content; and a significant improvement is made in the prediction of popular social contents that draw flash crowds.

We looked into ways to distribute network resources at edge CDN servers using social predictions based on data gathered from social networks, as well as the potential to infer users' preferences for contents based on their social profiles and behaviours. In a CPCDN, user social interactions and activities give the new design space the opportunity to investigate how to employ user intelligence for content distribution.

Reference

1. B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Web content cartography," in Proc. ACM Internet Meas. Conf., 2011, pp. 585–600.
2. E. Bakshy, J. Hofman, W. Mason, and D. Watts, "Everyone's an influencer: Quantifying influence on Twitter," in Proc. ACM Int. Conf. Web Search Data Mining, 2011, pp. 65–74.
3. B. Brewington and G. Cybenko, "How dynamic is the web?," Comput. Netw., vol. 33, no. 1, pp. 257–276, 2000.
4. M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon, "I tube, you tube, everybody tubes: Analyzing the world's largest user generated content video system," in Proc. ACM SIGCOMM, 2007, pp. 1–14.
5. M. Cha, A. Mislove, and K. Gummadi, "A measurement-driven analysis of information propagation in the Flickr social network," in Proc. ACM Int. Conf. World Wide Web, 2009, pp. 721–730.
6. F. Chen, K. Guo, J. Lin, and T. La Porta, "Intra-cloud lightning: Building CDNs in the cloud," in Proc. IEEE Int. Conf. Comput. Commun., Mar. 2012, pp. 433–441.
7. A. Datta, K. Dutta, H. Thomas, D. VanderMeer, and K. Ramamritham, "Proxy-based acceleration of dynamically generated content on the world wide web: An approach and implementation," ACM Trans. Data-base Syst., vol. 29, no. 2, pp. 403–443, 2004.
8. I. Dhillon, "Co-clustering documents, and words using bipartite spectral graph partitioning," in Proc. ACM SIGKDD Conf. Knowl. Discovery Data Mining, 2001, pp. 269–274.
9. A. Di Marco and R. Navigli, "Clustering and diversifying web search results with graph-based word sense induction," Comput. Linguistics, vol. 39, no. 3, pp. 709–754, 2013.
10. N. B. Ellison, C. Steinfield, and C. Lampe, "The benefits of Facebook friends: Social capital and college students' use of online social network sites," J. Comput.-Mediated Commun., vol. 12, no. 4, pp. 1143–1168, 2007.
11. B. Huberman, D. Romero, and F. Wu, "Social networks that matter: Twitter under the microscope," Comput. Res. Repository, vol. abs/0812.1045, 2008.
12. S. Kommula, I. Hsu, R. Jalan, and D. Cheung, "Global server load balancing," U.S. Patent 7,254,626, Aug. 7, 2007, et.al..

13. H. Kwak, C. Lee, H. Park, and S. Moon, “What is Twitter, a social network or a news media?,” in Proc. ACM Int. Conf. World Wide Web, 2010, pp. 591–600.
14. I. Lazar and W. Terrill, “Exploring content delivery networking,” IT Prof., vol. 3, no. 4, pp. 47–49, 2001.
15. T. Lemlouma and N. Layada, “Adapt ed content delivery for different contexts,” in Proc. IEEE Symp. Appl. Internet, Jan. 2003, pp. 190–197.