

PROACTIVITY IN CONTENT DELIVERY NETWORKS RESOURCE MANAGEMENT: THE STOCK OPTIONS CASE¹

Elias Vathias, Matthias Filippou Doukas, Panagiotis Giannopoulos and Stathes Hadjiefthymiades
Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens, Greece

NOTE: Corresponding author: Elias Vathias, evathias@di.uoa.gr

Abstract – This paper focuses on the management and pricing of resources in Content Delivery Networks (CDNs), which are extensively adopted nowadays as a very efficient mechanism for Internet information provision. We elaborate on a resource management model that is aligned with the concepts and mechanisms in capital markets. We use the concept of Stock Options (SOs) to address the scarcity and potential unavailability of CDN resources. Using a Predictive Reservation Scheme (PRS), network resources (traffic volume) are being monitored through kernel estimators in a given time frame. A Secondary Market (SM) significantly improves the efficiency and robustness of the PRS by allowing the fast exchange of unused resources (stocks) and SOs between the Origin Servers (OSs). This exchange can happen by allowing automatic electronic double auctions at the end of each day or at shorter time intervals. As a result, OSs may acquire resources (if needed) at standard prices, avoiding penalizing tariffs for last-minute requests. The efficiency of our prediction reservation scheme further improves.

Keywords – Content delivery networks, pricing plans, resource allocation and management, secondary market, stock options

1. INTRODUCTION

Nowadays, for many content-generating organizations, Content Delivery Networks (CDNs) are common practice. With content that becomes multimedia-richer, "heavier," and, structurally, much more complex, the Origins Servers' (OSs) population is continuously rising. We seek to optimize the reservation and usage of CDN resources from both the client (OS) and the CDN side.

Central to this design is the gradual shifting of the CDN's attention to more fine-grained resource management schemes that minimize overheads, such as unused capacity. The CDN, in turn, leases resources from network and storage providers (meta-CDNs). The traded resources are finite; hence, the CDN needs to plan for their acquisition and management carefully. Schemes and techniques for more accurate resource claims are of common interest and, in our domain, they positively impact the involved clients (OSs), the CDN operator/provider, and the meta-CDN. Improved resource utilization can benefit both sellers and buyers. CDN providers can deliver better services at more competitive prices to a much wider audience, especially in relation to other CDN providers that do not adopt the proposed framework. Furthermore,

improving resource utilization may help reduce network congestion.

In the previous studies detailed in [1]–[3] we adopt an optimized framework (Fig. 1) for the methodical management of CDN resources.

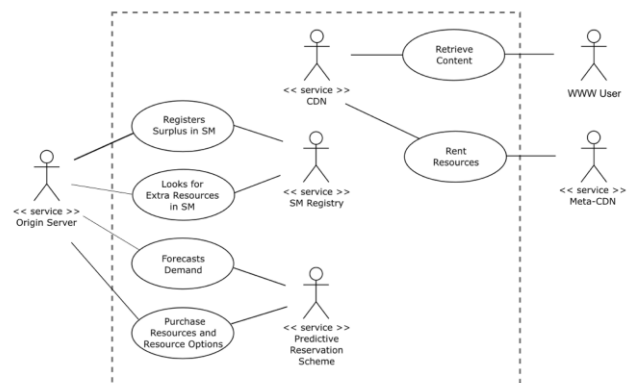


Fig. 1 - A resource market model for CDN: use case diagram

Specifically, we borrow concepts, schemes, and techniques from the capital market [4], [5]. We treat CDN resources as stocks. The stockholders are the content-generating organizations (OSs). The CDN monitors the incoming traffic for each client to better follow the exact resource use, protect against resource misuse, and establish load predictions for the immediate future. According to the client's needs for resources, stocks are purchased

¹ This paper extends the work of the authors published in [1].

dynamically to sustain the experienced visitor load efficiently. The CDN learns the temporal distribution of resource use for each client. This knowledge is of utmost importance to the rationalization of CDN resource management since the CDN can reserve resources for the OS well beforehand and cope with the fluctuating load.

OSs may experience loads incompatible with their expectations/predictions. Resources can be traded between clients to improve the resilience of the primary mechanism to prediction failures. A tool for such trading of CDN resources is the capital Secondary Market (SM).

The capital Stock Option (SO) is another tool already introduced that improves the resilience of forecasting and resource swapping mechanisms. CDN purchases SOs, for each OS, at prices and for the duration defined by models such as that of Black-Scholes (BS) [6], [7], Barone-Adesi & Whaley [8], Bjerksund & Stensland, and Ju & Zhong [9]. These SOs, are exercised if needed (i.e. if an OS runs out of resources before the end of the billing period and cannot find any available resources through the SM).

This paper reinforces the above mechanisms by introducing the ability to exchange unused SOs in the SM. A detailed description of the proposed mechanism and its modifications is provided in Section 3. We assess the performance of the proposed scheme using extensive anonymized real traces taken from high-traffic OSs such as governmental agencies, universities, popular sports content websites, and e-commerce websites.

Our findings show that exchanging unused SOs in the SM enables the proposed framework to further "absorb" prediction failures and optimize CDN resource management. We argue that this framework that borrows concepts, schemes, and techniques from the capital market is a good candidate for the modern Worldwide Web (WWW) ecosystem structure.

The paper is organized as follows. Section 2 discusses related work. Section 3 presents the architecture of the proposed framework and the modifications we introduce and evaluate in this article. Section 4 describes in detail the parameters of the simulations and the corresponding metrics and results. Finally, Section 5 concludes the paper and discusses future work.

2. RELATED WORK

Internet resource pricing is a crucial factor in the efficient allocation of Internet resources and the determinant of profit [10]. CDNs have to deal with the cost of interconnection and traffic on their networks. At the same time, they are trying to increase their revenues by choosing effective pricing strategies that have to do with the efficient allocation of Internet resources.

In view of the above, there are three primary pricing models: flat pricing, usage-based pricing, and congestion pricing.

Looking back at the early days of the Internet, users utilized few network resources. Thus, Internet Service Providers (ISPs) [11]–[13], aiming to attract many users, adopted a flat-rate billing, charging users based on access costs. Flat pricing was easy to implement and also stimulated network usage. However, drawbacks emerged as the increase in network content and the lack of incentives for efficient network resource usage resulted in increased traffic and overall network performance degradation. To address this, usage-based pricing [14]–[16] was proposed. The main idea was that if the charges were usage-based, and since Edell and Varaiya [17] showed that users are susceptible to pricing, a fairer and more efficient use of resources would be feasible. Among the problems that had to be addressed were the privacy issues in processing audits and the charge of non-expected traffic, such as advertisements and spam. For example, the 95th percentile pricing became an industry standard. In this pricing model, the peak flow within 5% of the total time (36 hours per month) is free of charge.

However, network traffic continued to increase, exacerbating congestion. As a result, the pricing mentioned above became more complex, leading to a relatively dynamic pricing model, "congestion pricing" [11], [18]–[25], which has been studied extensively. Congestion pricing dynamically sets prices that can reflect approximate real-time network resource usage and, especially when the network is busy, encourages shifting the traffic from peak time to non-peak time, reducing congestion.

Here it is worth noting that, in the CDN context, congestion reduction is not a key goal of pricing. Content providers subscribe to CDN services to overcome, among other things, Denial of Service (DoS) attacks or flash crowds. Moreover, the traffic of various content providers is unlikely to surge simultaneously. So CDNs can temporarily adjust

their infrastructure to handle the traffic spike and improve content availability. As a result, while relevant, the research on congestion pricing does not directly affect the CDN.

With regard to the pricing of the different models applied, pricing mechanisms can be categorized into two types: best-effort and Quality of Service (QoS) enabled. In the best-effort type, users are charged according to access rate or resource usage. In the QoS-enabled type, ISPs tend to serve different data streams with different QoS and price levels. Generally, for best-effort networks, pricing is always done at the edge of networks and incurs a lower overhead cost. On the other hand, QoS guaranteed services involve a higher audition and accounting cost. Priority-based pricing was first proposed by Cocchi et al. [26], [27] to perform service layering with the corresponding pricing. Another well-known QoS differentiation proposal was Odlyzko's Paris Metro Pricing [28], which divides the network into subnets and charges them differently. QoS guaranteed network architectures (e.g. IntServ and DiffServ) and their corresponding pricing mechanisms have been widely studied [29]–[34].

There are two main models for determining the appropriate price levels regarding network pricing methods: system optimization (models) and strategic optimization (models). System optimization models are mainly based on an optimization theory like the concept of Network Utility Maximization (NUM) framework proposed by Kelly [35], which is the initial work of Internet system optimization, as well as other works on network utility maximization [36]. Strategic optimization models, i.e. considering the strategic behaviors of others when setting prices or making other decisions [37], [38], are based on non-cooperative games [39], [40].

Following those previously mentioned, several recent studies have been conducted on various aspects of caching and CDN technologies, including resource management and pricing.

2.1 Caching

In the area of **multilevel caching**, it is found that certain clients exhibit "aggressive" behavior, frequently monopolizing the cache disk space. As a result, other clients are confined to the remaining disk space, with the consequence of expelling "important" objects, thus experiencing low performance due to many cache misses.

The authors in [41] proposed a framework that discourages monopolizing the cache disk space by a minority of clients while rewarding clients who contribute to the overall hit rate by offering them more disk space. Hosanagar et al. [42], [43] argue that the use of best-effort caching will decrease as OSs move towards dynamic content while seeking accurate business intelligence regarding website usage. They argue that CDNs can play an essential role in intermediating between OSs that seek the benefits of edge delivery and ISPs that can install servers at the network's edge.

2.2 Distributed group of nodes

Adopting a **distributed group of nodes** is another standard model for studying caching proxies, CDN, and Peer to Peer (P2P) technologies. In such a model, each node uses the storage capacity to create copies of objects and render them available to local and remote users. The creation of copies of objects can be done through replication (permanent copies) or caching (temporary copies). In the first case, the authors in [44] propose a Two Step Local Search (k) algorithm that protects nodes from mismanagement. In the latter case, the authors in [45] propose detecting, addressing, and adjusting mismanagement mechanisms.

2.3 CDNs

In the **CDN area**, when several Service Classes (SCs) with different QoS are available, the authors in [46] discussed a simple differentiated service type architecture to provide fair service to the subscribed publishers. In this architecture, the authors determine the optimal pricing, the optimal allocation of resources, and the optimal number of services to be offered.

In similar research [47], using analytical models, the authors addressed the optimal pricing of the offered services and studied how external factors such as the cost of bandwidth and security issues affect pricing. For example, they found that (a) declining bandwidth costs will negatively impact CDN revenues and profits, (b) growing content distribution-related security concerns will lead CDNs to invest in technology to alleviate these problems or reduce prices, and (c) more extensive CDN networks can charge higher prices in equilibrium which should strengthen their economies of scale and make it more difficult for entrants to compete against them.

Hosanagar et al. [48] study the monopoly CDN regarding optimal pricing. They argue that

traditional usage-based pricing plans should offer volume discounts unless subscribing content providers have highly heterogeneous traffic flows. Moreover, they claim that percentile-based pricing plan profitability can be substantially higher than traditional usage-based billing.

The authors in [49] investigated the maximization of the benefits gained by the OS and CDN. Among the results of their research, they found that (a) the reduction in the price offered by the CDN results in higher investment by the OSs but lowering the price more than a specific level results in the revenue reduction since the CDN has a limited cache space and the OSs' requests cannot be completely satisfied (b) the surrogate revenue is maximized when the CDN cache space is equal to the total publisher demand, (c) surrogates should not be very close to the users and (d) during an optimum system investment while the total of publisher utilities are maximized, some of the publisher's individual utilities are reduced.

2.3.1 Current CDN technology

Besides "classic" single CDNs, various other solutions are offered to the OSs: **Cloud CDNs (CCDNs)** combine resources (compute, storage, and network that are available as web services) with CDN services. Broberg et al. [50] introduced a system that uses storage cloud resources and deploys as many surrogates in user-requested locations as their storage and transfer budget allows, keeping them active as long as sufficient budget remains. Then they measure the utility of content delivery via this system and use it to devise a request-redirection policy that ensures high-performance content delivery [51].

Multi-CDN solutions are viewed as an overlay of existing, individual CDNs, offering more Points of Presence (POPs), optimal reach, and redundancy, which are essential when you want to reach an international or worldwide audience. For example, the authors in [52] refer to Netflix, which employs a blend of data centers and CDNs for content distribution and propose a solution that can improve the average bandwidth by more than 50%.

Often, OSs deliver their content to Multi-CDN through **content brokers**. Mukerjee et al. argue [53] that brokers have been shown to invalidate many traditional delivery assumptions (e.g. shifting traffic invalidates short and long-term traffic prediction) by not communicating their decisions to CDNs, which can have unintended consequences,

including higher costs. Then they analyze [54] these problems, examine the design space of potential solutions, and find that a marketplace design (inspired by advertising exchanges) potentially provides interesting trade-offs.

Private CDNs (P-CDNs) are used by companies with high bandwidth and throughput needs or in-house network capacity (such as telecommunication companies), or high security and compliance requirements (such as payment or government companies). P-CDNs do not share resources with other customers. Canali et al. [55] present P-CDN deployment by a company with a networking infrastructure that is outsourced to a third party. Their findings suggest that:

- The introduction of a P-CDN can increase by more than eight times the number of clients that can access the multimedia data at the highest quality;
- There is a significant trade-off between the performance of the P-CDN (in terms of media quality for the users) and the number of edge servers deployed; and
- As the cost of edge servers may be high for a single content provider, the parameter for deciding if a company branch is large enough to be chosen to host an edge server remains a critical decision for the tuning of the model.

Telco CDNs are telecommunication companies or Telecommunication Service Providers (TSPs) licensing CDN platforms or merging with existing CDN companies and taking advantage of lower bandwidth costs. However, owning a telecommunication network alone is not enough for a telco CDN to enjoy market benefits. Lee et al. [56] analyze the strategic interactions between CDNs and telco CDNs and study the conditions that can lead to alliances among telco CDNs. They also provide evidence that market benefits are possible if a telco CDN appropriately manages to offer better service quality by exploiting its competitive advantages (e.g. joint traffic engineering and content distribution). The authors in [57] show that there are cases under which the potential for complete resource pooling and revenue sharing among a telco CDN federation is beneficial, although in most cases, resource pooling on its own brings more benefits to each individual telecom operator. Frangoudis et al. [58] propose an architecture for on-demand service deployment over a telco CDN, where the CDN resources are leased dynamically in different regions based on customer demand.

Spagna et al. [59] offer design considerations for building telco CDNs with a slight focus on mobile networks. They discuss cache placement, request routing, and content outsourcing.

In **virtual (or virtualized) CDNs (vCDNs)**, virtual caches are deployed dynamically (as virtual machines or containers) in physical servers distributed across the provider's geographical coverage. A vCDN can be more cost-effective than a CDN running on dedicated infrastructure since virtual machines better utilize server resources. Frangoudis et al. [60] propose an architecture for **CDN as a Service (CDNaaS)**, allowing content providers to order and deploy the vCDN surrogate servers in ISPs. ISPs leverage CDNaaS to receive the content provider requests, orchestrate the resources, and deploy the surrogate server functionality on available infrastructure.

P2P CDNs are based on P2P technology and offer lower cost and high quality services by leveraging the massive fragmentation of idle resources in edge networks. Xu et al. [61] propose and analyze a novel hybrid architecture that integrates CDN and P2P-based streaming media distribution and significantly lowers CDN capacity reservation costs without compromising the media quality delivered. Yin et al. [62] evaluate the performance of LiveSky, a commercially deployed hybrid CDN-P2P live-streaming system, using data from these real-world deployments and argue that such a hybrid CDN-P2P system (a) provides quality and user performance comparable to a CDN and (b) effectively scales the system capacity when the user volume exceeds the CDN capacity.

2.4 Auctions in resource management

Regarding auctions in the resource management area, the authors in [63] propose an auction approach to dynamically allocate the spectrum in an SM to use the wireless spectrum better. The authors in [64] propose 'Progressive Second Price' auctions to allocate network bandwidth based on the demand and willingness to pay from competing users who decide their bids based on their accurate valuation. The authors in [65] consider a two-level hierarchical business model for selling bandwidth. In this model, a single vendor allocates bandwidth to intermediary service providers, who in turn sell their assigned shares of bandwidth to their customers at the lower level.

2.5 Our contribution

Our work differs from what has been studied and proposed in the related work section in the sense that (a) it uses load monitoring, modeling, and prediction to reserve the necessary resources proactively; (b) it uses the logic of capital market instruments (stocks, SOs, SM and electronic auctions), based on corporate finance to manage resources and automate their financial management and exchange; and (c) it offers an overall more fine-grained resource management scheme that reduces overheads such as unused capacity and enables the CDN to plan its resources' needs carefully. As a result, all the involved entities benefit from the rationalized use of CDN resources. Furthermore, improving resource utilization may help reduce network congestion.

The novelty of this paper, in relation to the previous studies [1]–[3], is (a) the introduction of the ability to exchange the available SOs between OSs through the SM, and (b) the investigation and understanding of the performance of the proposed solution by conducting simulations with different combinations of the primary metrics of the SOs (number and duration).

3. ARCHITECTURE

This section recapitulates the proposed framework's architecture so far, detailed in [1]–[3]. We then describe the SM's SOs exchange and the way it is integrated into the existing architecture. The model formulation of the proposed framework can be found in Appendix A1. In general, we have chosen the (proposed) market ecosystem to operate on a socially optimal concept of allocating resources among market participants without greedy or competitive behaviors. Our goal is to maximize the benefits for all the entities involved.

3.1 Existing architecture

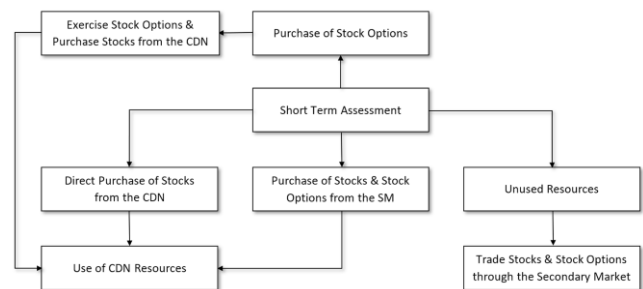


Fig. 2 - PRS

Overall, we adopt a PRS that involves four distinct (operational) aspects (Fig. 2). Such a scheme

attempts to accurately model the load that corresponds to each OS. This scheme uses such information to direct the appropriate forecasting mechanisms for immediate reservations of CDN resources. For the management of these resources, we adopt the financial instrument of stocks and SOs (rights to purchase CDN resources at a predetermined price and for a predetermined period, defined by models such as that of BS).

While clients seek the most accurate forecasts of the forthcoming load to prevent high or low reservations, predictions cannot be 100% accurate. OSs can exchange resources (stocks) via the SM at mutually favorable prices to minimize the predictive model's deficiencies. OSs still deficient in resources will eventually exercise the SOs they may have.

3.1.1 Load prediction mechanism

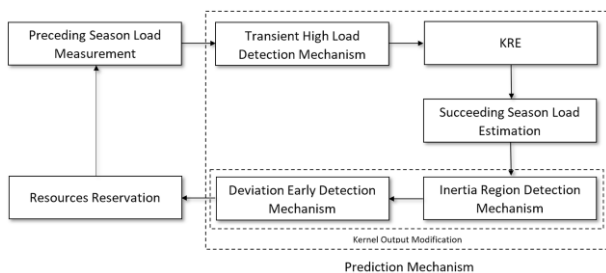


Fig. 3 – Load prediction mechanism

The Load Prediction Mechanism (LPM) itself (Fig. 3) is the first aspect of PRS [1]–[3] and is based on Kernel Regression Estimators (KREs), which continuously and accurately model the site workload over specific time frames and output a "refined" model of the anticipated load. The LPM is also based on other complementary techniques, such as:

- **Transient High Load Detection Mechanism (THLDM)**, which detects transient high load and excludes the corresponding load from the next season's forecasting;
- **Inertia Region Detection Mechanism (IRDM)**, which modifies the KRE output whenever the modeled load leaves local maxima;
- **Deviation Early Detection Mechanism (DEDM)**, which monitors the efficiency of the LPM regarding the actual load and adjusts the resources reserved for the next time frame; and
- **Initial Resource Reservation Monitoring Mechanism (IRRMM)**, for the management of the initial service period where no previous data exists.

The outcome of the LPM is binding as an upper limit for each OS's next resource reservation. This limitation protects the entire system by prohibiting malicious users (OSs) from reserving many resources in advance at low prices and selling them at a certain profit through the SM. More details on the LPM's use and its components can be found in Section 3.1 of [3].

3.1.2 Stocks

We adopt the financial instrument of stocks to implement the predictive resource reservation. In the context of our problem, stocks are bandwidth share units traded at specific prices.

An OS (the CDN client) buys stocks to immediately reserve resources at the currently available prices (Fig. 2). The OS essentially acquires the right to use the corresponding resources (bandwidth) for a specific duration and at a certain price by buying stocks. Such rights can be easily exchanged via the SM (figures 1 and 2).

The assumptions upon which our stock-based modeling relies are as follows: (a) The CDN resources are not infinite, i.e. the number of stocks that the CDN can trade at any price is countable, and (b) the CDN charging for resource reservation/use depends on the volume/quantity and the planned date/time of the expected use. Reservations that are confirmed early are preferentially priced. Urgent resource requests are "penalized" with high prices.

3.1.3 SM

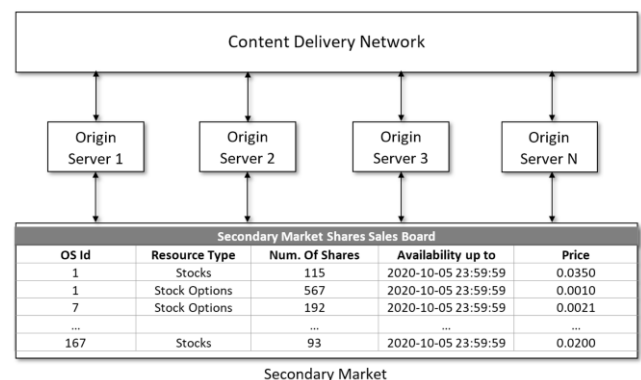


Fig. 4 - SM

The first line of defense for the predictive model's failures is the SM (figures 1, 2, and 4), where new financial instruments are not issued, nor are the issuers, such as corporations, raising new funds. Only investors can buy from other investors [5].

If the reserved resources of an OS are significantly underused, the OS can sell a percentage of them to

another OS via SM. Thus, the site reduces or eliminates unused yet reserved resources (balancing the prediction failure) and improves resource usage efficiency. Conversely, if the reserved resources of an OS are found to be insufficient, the OS can seek additional resources via the SM. Thus, the OS avoids being penalized (overcharged) by the CDN for the needed resources (as an implication of prediction failures) and manages to "absorb" the unforeseen load through the SM traded resources. Offers and sales in the SM take place through double auctions.

The SM's adoption has proven beneficial to all parties involved (buyer, seller, and CDN). The buyer adopts a cost-efficient scheme for unforeseen load, the seller reduces its potential loss (or even makes a profit), while the CDN increases its resource utilization percentage. The last enables the CDN to offer lower prices to its customers and improve its position in the competition. More details on the SM can be found in Section 3.3 of [3].

3.1.4 SOs

Following the SM's use, clients may still lack resources at the End of the Day (EoD) and have to turn to the CDN to buy resources at penalty prices. To avoid such an action, the LPM records the shortage of resources, makes forecasts, and purchases SOs at the beginning of each day (BoD) for each OS. SOs can be exercised if needed, and the OS can acquire resources at a standard price.

SOs act as the second line of defense against predictive model failures. More details on using the SOs in the proposed framework can be found in Section 3.4 of [1].

3.2 Architecture proposed modifications

As mentioned earlier, in this paper, we enhance the proposed framework mainly (a) by introducing the ability to exchange unused SOs in the SM and (b) by looking for the optimal combination of basic SO parameters (quantity and duration) that will result in the maximum possible reduction of stocks acquired at a penalty price and their replacement by stocks acquired through the exercise of the SOs. With this enhancement, we aim to improve the overall performance further.

Description of the overall function

Let us first recall one key feature of our model: The CDN serves all OS clients during the day, even if they run out of resources. At the EoD, any accounting issues are settled.

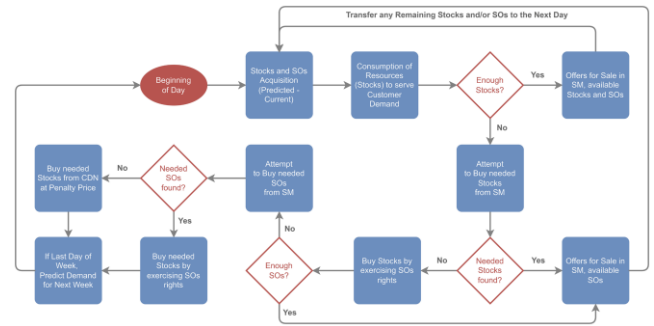


Fig. 5 - Overall function workflow

In the architectural version proposed in this paper (Fig. 5), at the beginning of each week and for each OS, the CDN acquires resource consumption forecasts (that further determine stocks' and SOs' needs) for the following seven days. Moreover, at the BoD, each OS buys a) the number of stocks that are forecasted minus the number of stocks that may remain unused from the previous day, and b) the number of SOs that are forecasted (a percentage of forecasted stocks in this proposed architecture modification) minus the number of active SOs that may remain unused from the previous days. Then and by the EoD, each OS faces one of the following scenarios:

Scenario 1: The stocks available to the OS are sufficient to cover its daily consumption. Consequently, no options are exercised. At the EoD, the OS participates in the SM auctions as a seller, offering (a) all remaining stocks at their purchasing price and (b) from the available SOs, only those that expire at the closest date at a price depending on their number and expiry date. Any unsold stocks and SOs remain available for the next day.

Scenario 2: The stocks available to the OS run out before the EoD. The OS participates in the SM auctions as a buyer, seeking to buy stocks at lower prices than the penalized price set by the CDN.

- If the stocks (that the OS may find in the SM) are sufficient to meet its current (day) needs, then the OS participates in the SM auctions as a seller offering, from the available SOs, only those that expire at the closest date at a price depending on their number and expiry date. Any unsold SOs remain available for the next day.
- If the stocks are insufficient, the OS exercises its SOs and buys stocks at a predetermined price, preferring those that expire sooner.
 - Subsequently, if the available SOs are sufficient to meet its needs, then the OS participates in the SM auctions as a seller offering, from the possibly remaining SOs,

only those that expire at the closest date, again at a price depending on their number and expiry date. Those not sold remain available for the following day.

- If the available SOs are insufficient, then, after exercising them all, the OS participates in the SM auctions as a buyer, seeking to buy SOs, at prices again set by the seller. Finally, if even these SOs (that the OS may find in the SM) are insufficient to cover its daily consumption, it buys the necessary stocks from the CDN at a penalty price.

3.3 Framework implementation

The proposed framework can be implemented as software that runs on the accounting component of the CDN, evaluates the recorded information from the previous season, estimates future workload for each OS, and enables the trading of resources and SOs between OS through the SM. Matlab, C, and Java code were used for the simulation process. The total computation time (on a system with a quad processor and 8 GB RAM) for emulating the operation of six OSs, one CDN, and one meta-CDN, for 52 weeks (1 year) is approximately 36 minutes.

4. SIMULATION AND RESULTS

In the previous studies [1]–[3], we compared and evaluated various pricing plans (variants) of the proposed framework with other pricing plans available on the market. In this paper, we compare pricing plan 6a of the proposed framework (which includes the use of SOs) with a new variant of the proposed framework, namely pricing plan 7a, which is based on plan 6a and also allows the exchange of SOs to the SM. We further process and analyze these two variants, varying the duration and number of SOs.

Other pricing and resource management models used in the broader context of networks and distributed systems operate differently from the proposed framework, do not directly relate to the operation of CDNs and were therefore not evaluated.

Please note that, during these simulations, we consider that the costs to be reported below include all other CDN charges to OSs, such as management fees and request and ingress costs.

The rest of this section is structured as follows: Section 4.1 describes the main parameters of the simulations we perform. Section 4.2 describes the metrics defined as the essential elements of observation of the simulation, while in Section 4.3,

we define the scenarios that we choose to simulate and compare. In Section 4.4, we report on our findings.

4.1 Simulation parameters

The simulations involve 6 OSs, 1 CDN, and 1 meta-CDN, are trace-driven, based on anonymized Apache web server cache logs of primary OSs, and cover a duration of 364 days (52 weeks).

More specifically, the logs have been obtained from the following content providers: two chambers of commerce portals (with business information and personalized services), a university website (with student information and personalized course and curriculum services), a popular sports content website (with details on race dates and results, personal drivers pages, national records, and statistics) as well as two popular e-commerce websites.

Table 1 summarizes the characteristics of the employed real traces and the general simulation parameters.

Table 1 - General simulation parameters

| Characteristics | Values |
|--------------------------------|-----------------------------|
| OSs | 6 |
| Duration of Simulation Period | 364 days |
| Forecasting Reference Period | One week |
| KRE Kernel Function | Gaussian |
| KRE Kernel Bandwidth | Varying from 0.8 to 5.1 |
| Control Time Granularity | 24 hours |
| Min Volume Penalty Calculation | 1 byte |
| Service Classes | 6 |
| Penalty Type | High (120% of Normal Price) |

Note: Control time granularity indicates the rate at which measurements are taken, which is, in turn, used for the forecasting task.

Regarding the SOs, the parameters depicted in Table 2 have been adopted across all the investigated scenarios. SOs can be exercised at any time, including their expiration (maturity) date and not just their expiration date. The Black-Scholes formula is used for the SO price calculation as the SOs do not pay dividends [6], [7].

Table 2 - SO-specific simulation parameters

| Parameter | Value |
|------------------------------|--------------------------|
| Annual Risk-free rate | 5.00 % |
| Annualized Volatility | 1.00 % |
| Dividend (Continuous Method) | 0.00 |
| Yield Rate (SO) | 0.00 |
| Time to Maturity | 7 or 14 or 21 or 28 days |

Table 3 shows the standard and penalty price of stocks and the purchase price of SOs per Traffic Volume (TrV) and SC, paid by the OS. More information on how the values of Table 3 were calculated can be found in sections 3.4.1, 3.4.2, 4.2.4, and 4.2.6 of [1].

Table 3 - Costs and TrV details per service class

| SC | TrV (GB per day) | Stocks | | SOs (Time to Maturity) | | | |
|----|---------------------|-------------------------|------------------------|------------------------|------------------------|------------------------|------------------------|
| | | Standard (\$ per GB) | Penalty (\$ per GB) | 7 days (\$ per GB) | 14 days (\$ per GB) | 21 days (\$ per GB) | 28 days (\$ per GB) |
| P1 | 0 - 0.357 | 0.035 | 0.042 | 0.002031629 | 0.002063227 | 0.002094795 | 0.002126333 |
| P2 | >0.357 - 3.57 | 0.030 | 0.036 | 0.002026836 | 0.002053647 | 0.002080432 | 0.002107192 |
| P3 | >3.57 - 35.71 | 0.025 | 0.030 | 0.002022044 | 0.002044067 | 0.002066069 | 0.002088050 |
| P4 | >35.71 - 357.14 | 0.020 | 0.024 | 0.001018210 | 0.001036403 | 0.001054579 | 0.001072737 |
| P5 | >357.14 - 3,571.42 | 0.015 | 0.018 | 0.001013418 | 0.001026824 | 0.001040216 | 0.001053596 |
| P6 | >3,571.42 | 0.010 | - | 0.001010543 | 0.001021076 | 0.001031598 | 0.001042111 |

4.2 Metrics

The following metrics have been considered to compare the effectiveness of the different scenarios of the proposed framework.

- **Predicted traffic:** The predicted traffic by our LPM.
- **Stocks acquired at the BoD per OS and scenario:** The stocks (GB) acquired at standard prices, at the BoD, and their respective costs (\$).
- **Stocks exchanged between OSs, via SM per OS and scenario:** The stocks (GB) exchanged, during each day, between OSs via the SM and their respective costs (\$).
- **Stocks acquired by exercising SOs, per scenario:** The stocks (GB) acquired, during each day, by exercising SOs and their respective costs (\$).
- **Stocks acquired at penalty prices per OS and scenario:** The stocks (GB) acquired, during each day, at penalty prices from CDN and their respective costs (\$).
- **SOs acquired at the BoD per OS and scenario:** The SOs (GB) acquired at the BoD and their respective costs (\$).
- **SOs exchanged between OSs via SM per OS and scenario:** The SOs (GB) exchanged, during each day, between OSs via the SM and their respective costs (\$).
- **SOs exercised per OS and scenario:** The SOs (GB) exercised each day and their respective costs (\$).
- **Unused SOs' cost per scenario and OS:** The SOs (GB) that were not exercised until their expiration day and their respective costs (\$).

4.3 Comparative assessment

We simulate the operation and compare the efficiency of various scenarios determined by the following parameters of the model:

- **SOs exchange in SM:** True / False (This parameter determines whether the SOs are traded on the SM during the simulation. Please note that resources are exchanged in the SM, regardless of whether the SOs are traded).
- **Maximum SOs duration (in days):** 6 / 7 / 14 / 21 / 28 (This parameter defines the validity period of the SOs that can be exercised during the simulation).
- **Percentage of SO purchased at the BoD based on the initially predicted traffic:** 5% / 10% / 15% (The percentage of SOs to be purchased relative to the predicted resources to be purchased by each OS at the BoD).

Table 4 - Simulated scenarios

| Scn. | Pricing Plan | SOs Exchange in SM | Max SO Duration (days) | Percentage of SO Purchased |
|------|--------------|--------------------|------------------------|----------------------------|
| 1 | 6a | False | 7 | 10% |
| 2 | 6a | False | 14 | 10% |
| 3 | 6a | False | 21 | 10% |
| 4 | 6a | False | 28 | 10% |
| 5 | 7a | True | 7 | 10% |
| 6 | 7a | True | 14 | 10% |
| 7 | 7a | True | 21 | 10% |
| 8 | 7a | True | 28 | 10% |
| 9 | 6a | False | 6 | 15% |
| 10 | 7a | True | 6 | 15% |
| 11 | 6a | False | 14 | 5% |
| 12 | 7a | True | 14 | 5% |

Table 4 presents the number of scenarios tested along with the main simulation parameters. In total, twelve (12) scenarios have been chosen to be simulated through the proposed framework.

4.4 Results

This section presents the simulation results, following the metrics we described in Section 4.2. We show that the ability to exchange SOs in the SM results in the reduction of stocks that the OSs acquire at penalty prices and thus the cost they pay. In addition, we show that a further increase, beyond the 15% percentage, in the number of SOs purchased from each OS at the BoD is needed, with a duration of 14 days or more. These findings will be further analyzed in the following subsections.

While presenting the following and in order to compare the individual results, we use the concept of Percentage Difference (PD), which we define as

follows:

$$PD(x, y) = \frac{|x-y|}{\frac{x+y}{2}} * 100 \quad (1)$$

In summary, this section includes the following:

- In Section 4.4.1, we refer to predicted traffic.
- In sections 4.4.2 to 4.4.5, we detail the results related to the acquisition of stocks in four distinct situations: (a) at the BoD based on predictions, (b) via exchange in the SM, (c) by exercising SOs, and (d) at penalty prices.
- In Section 4.4.6, we refer to the proposed mechanism’s new version improvement, which gives the possibility to search for SOs in the SM; the reduction in acquisition of stocks at penalty prices, and the replacement of them with stocks acquired through the exercise of more SOs, precisely because of the possibility of searching for them in SM.
- Section 4.4.8 summarizes the acquisition of stocks and their cost, while Section 4.4.7 describes detailed statistics on the acquisition and use of SOs.

Thus, one can focus on sections 4.4.6, 4.4.7, and 4.4.8, while for more details, one can also refer to sections 4.4.2 - 4.4.5.

4.4.1 Predicted traffic

We start the presentation regarding the predicted traffic for each of the six OSs for the entire simulation period (Table 5). The wide range of TrV helps us better study the plans' efficiency and behavior under consideration. The predicted traffic is a metric unaltered by each scenario's different parameters, as it depends solely on the LPM and is not affected by the actions that follow.

Table 5 - Predicted traffic (GB)

| | A | C | E | J | L | P | Total |
|-------------------------------|------------|-----------|----------|-------|-----------|--------|-------------------|
| Predicted Traffic (GB) | 114,712.53 | 18,405.69 | 4,223.55 | 88.70 | 15,190.37 | 478.10 | 153,098.94 |

4.4.2 Stocks and their cost, acquired at the BoD per OS and scenario

The stocks (GB) acquired at the BoD and the cost (\$) of acquiring them (Table 6) remain undifferentiated in all the scenarios under consideration due to the order of activation of the proposed architecture's mechanisms. The column "Total" also reflects the total stocks sold by the CDN at the BoD and the corresponding revenue.

Table 6 - Stocks acquired at the beginning of day per OS and scenario

| | OS 1 | OS 2 | OS 3 | OS 4 | OS 5 | OS 6 | Total |
|--------------------|------------|-----------|----------|-------|-----------|--------|-------------------|
| Stocks (GB) | 105,728.57 | 15,786.36 | 3,860.41 | 79.29 | 12,423.82 | 414.92 | 138,293.36 |
| Cost (\$) | 3,044.54 | 369.24 | 96.53 | 2.75 | 303.33 | 12.42 | 3,828.80 |

4.4.3 Stocks and their cost, exchanged between OSs, via SM per OS and scenario

The stocks (GB) exchanged via SM each day, and the cost (\$) of this exchange (Table 7) remain undifferentiated in all the scenarios under consideration due to the order of activation of the proposed architecture's mechanisms. The CDN is not involved in this process, as the stocks are exchanged between the OSs.

Table 7 - Stocks exchanged between OSs, via SM per OS and scenario

| | OS 1 | OS 2 | OS 3 | OS 4 | OS 5 | OS 6 | Total |
|---------------------------------------|-----------|-----------|---------|-------|-----------|--------|-----------|
| Stocks Acquired via SM (GB) | 2,356.60 | 1,752.52 | 467.48 | 8.76 | 1,829.57 | 80.37 | 6,495.29 |
| Stocks Sold via SM (GB) | -2,883.90 | -1,807.52 | -340.34 | -5.29 | -1,410.30 | -47.95 | -6,495.29 |
| Cost of Stocks Acquired via SM (\$) | 56.86 | 44.51 | 10.73 | 0.24 | 44.50 | 2.09 | 158.92 |
| Revenue from Stocks Sales via SM (\$) | -74.08 | -41.48 | -8.10 | -0.17 | -33.69 | -1.40 | -158.92 |

Here we observe the following trend: OSs that receive the most traffic (and therefore potentially acquire stocks at lower prices) sell more stocks through the SM than they buy. On the contrary, OSs that receive less traffic (and acquire stocks at higher prices) acquire more stocks through the SM than they sell.

4.4.4 Stocks and their cost, acquired by exercising SOs, per OS and scenario

This metric refers to the stocks acquired by exercising SOs (Table 8) regardless of whether these SOs were acquired at the BoD or through the SM. More information on buying, exchanging, and exercising SOs can be found in Section 4.4.7.

Table 8 - Stocks acquired by exercising SOs, per OS and scenario

| Scn. ID | OS 1 | OS 2 | OS 3 | OS 4 | OS 5 | OS 6 | Total |
|---------|----------|----------|--------|------|----------|-------|----------|
| 1 | 2,714.14 | 319.41 | 42.67 | 0.74 | 404.44 | 7.39 | 3,488.80 |
| 2 | 2,728.67 | 355.18 | 54.72 | 0.90 | 457.27 | 8.09 | 3,604.83 |
| 3 | 2,753.58 | 371.26 | 60.48 | 0.99 | 462.57 | 8.58 | 3,657.46 |
| 4 | 2,760.36 | 370.93 | 60.45 | 1.00 | 465.75 | 8.78 | 3,667.26 |
| 5 | 2,960.46 | 725.31 | 134.65 | 3.40 | 1,220.20 | 12.81 | 5,056.83 |
| 6 | 3,064.09 | 715.76 | 137.30 | 3.29 | 1,255.52 | 13.24 | 5,189.20 |
| 7 | 2,967.53 | 828.98 | 136.01 | 3.67 | 1,260.70 | 13.77 | 5,210.66 |
| 8 | 2,967.68 | 830.02 | 136.01 | 3.67 | 1,260.70 | 13.77 | 5,211.85 |
| 9 | 3,254.51 | 455.64 | 52.89 | 0.59 | 563.59 | 9.04 | 4,336.26 |
| 10 | 3,624.74 | 1,091.98 | 172.96 | 4.21 | 1,548.40 | 16.32 | 6,458.61 |
| 11 | 1,653.53 | 196.85 | 33.08 | 0.63 | 236.96 | 4.64 | 2,125.68 |
| 12 | 1,791.44 | 467.58 | 80.43 | 2.32 | 597.72 | 7.71 | 2,947.21 |

By carefully observing the results of the simulations in Table 8, we conclude the following:

- Regarding the duration of the SOs:** From scenarios 1-4, we conclude that, with the SM inactive for the exchange of SOs, increasing the SOs duration from 7 to 14 days results in an increase of approximately 3.27% (percentage difference) in the total acquired stocks (through SOs). A further increase in SOs duration results in a slight additional increase (1.45% from 14 to 21 days and 0.27% from 21 to 28 days) in the total stocks acquired (through SOs). From scenarios 5-8, we conclude that, with the SM active for the exchange of SOs, an SOs duration increase from 7 to 14 days results in an increase of approximately 2.58% (percentage difference) in the total acquired stocks (through SOs). A further increase in the SOs duration results in a slight additional increase (0.41% from 14 to 21 days and 0.02% from 21 to 28 days) in the total stocks acquired. Thus, we can conclude that the acquisition of stocks through the exercise of SOs, increases with the increase of the latter's duration from 7 to 14 days. A further increase in duration to 21 or 28-days implies smaller increases in the acquisition of stocks.
- Regarding the SM's activation for the SOs exchange:** We conclude that the activation of SM for the trade of SOs (SM is always on for exchanging stocks) results in a significant increase in the acquired stocks through the exercise of SOs. More specifically, for an SOs duration of 7 days (scenarios 1 and 5), there is an increase of 37.70% in the acquired resources (through SOs); for 14 days (scenarios 2 and 6 and 11 and 12), there is an increase of 36.03% and 32.39% respectively; for 21 days (scenarios 3 and 7) there is an increase of 35.03%; for 28 days (scenarios 4 and 8) there is an increase of

34.79%, and for 6 days (scenarios 9 and 10) there is an increase of 39.32%.

- Regarding the percentage of SOs acquired:** We conclude that an increase in the percentage of SOs acquired results in a significant increase in the acquired resources through SOs, whether or not the SM is activated. More specifically, for an increase of 5% in the acquired SOs, we get the following results: for scenarios 11 and 2 there is a 51.62% increase, for scenarios 12 and 6 a 55.11% increase, for scenarios 5 and 10 a 24.35% increase and for scenarios 1 and 9 a 21.66% increase.

Finally, we conclude that between the SOs duration and the percentage of SOs acquired at the BoD, the latter affects more the stocks acquired through SOs (scenarios 1,2,3,4,9 and 11, and 5,6,7,8,10 and 12).

Regarding the same process, i.e. that of acquiring stocks through OSs' exercise, Table 9 indicates the relevant costs.

Table 9 - Cost of stocks acquired by exercising SOs, per OS and scenario

| Scn. ID | OS 1 | OS 2 | OS 3 | OS 4 | OS 5 | OS 6 | Total |
|---------|-------|-------|------|------|-------|------|--------|
| 1 | 58.36 | 7.85 | 1.22 | 0.02 | 9.88 | 0.24 | 77.58 |
| 2 | 58.40 | 8.75 | 1.55 | 0.03 | 11.08 | 0.27 | 80.07 |
| 3 | 58.75 | 9.13 | 1.71 | 0.03 | 11.19 | 0.28 | 81.09 |
| 4 | 58.78 | 9.12 | 1.71 | 0.03 | 11.25 | 0.29 | 81.17 |
| 5 | 64.86 | 17.56 | 3.55 | 0.10 | 27.99 | 0.41 | 114.48 |
| 6 | 66.75 | 16.97 | 3.62 | 0.10 | 28.74 | 0.42 | 116.62 |
| 7 | 64.54 | 19.65 | 3.59 | 0.11 | 28.87 | 0.44 | 117.20 |
| 8 | 64.54 | 19.68 | 3.59 | 0.11 | 28.87 | 0.44 | 117.23 |
| 9 | 67.06 | 10.81 | 1.49 | 0.02 | 13.27 | 0.30 | 92.94 |
| 10 | 75.12 | 25.49 | 4.55 | 0.13 | 34.49 | 0.52 | 140.31 |
| 11 | 38.11 | 5.36 | 0.95 | 0.02 | 6.09 | 0.15 | 50.68 |
| 12 | 41.89 | 11.91 | 2.13 | 0.07 | 14.78 | 0.25 | 71.04 |

By carefully observing the results related to stock acquisition costs (Table 9), we conclude that the same observations hold as those of the acquired stocks (Table 8). The columns "Total" in tables 8 and 9 also reflect the total stocks sold by the CDN to the OSs through exercising SOs and the corresponding revenue.




Fig. 6 - Resources acquired via SOs per OS and scenario

The previous analysis is depicted in Fig. 6, where we compare the corresponding scenarios for the stocks acquired and the acquisition costs. The only difference between the comparable scenarios is whether or not the SOs are traded on the SM.

4.4.5 Stocks and their cost, acquired at penalty prices per OS and scenario

Following the previous metric, we are considering the number of resources purchased at penalty prices since these costs significantly affect the OS's wallet.

Table 10 - Stocks acquired at penalty prices per scenario and OS (GB)

| Scn. Id | 1 | 2 | 3 | 4 | 5 | 6 | Total |
|---------|----------|----------|--------|------|----------|-------|-----------|
| 1 | 6,263.96 | 2,578.86 | 183.01 | 4.80 | 2,338.36 | 22.95 | 11,391.94 |
| 2 | 6,222.32 | 2,554.37 | 178.43 | 4.72 | 2,293.73 | 22.34 | 11,275.91 |
| 3 | 6,195.90 | 2,540.10 | 174.79 | 4.65 | 2,285.89 | 21.95 | 11,223.28 |
| 4 | 6,192.18 | 2,539.23 | 174.69 | 4.63 | 2,281.01 | 21.74 | 11,213.48 |
| 5 | 6,037.01 | 2,160.26 | 100.73 | 2.23 | 1,506.04 | 17.63 | 9,823.91 |
| 6 | 6,030.73 | 2,069.71 | 100.32 | 1.97 | 1,472.09 | 16.71 | 9,691.54 |
| 7 | 6,020.29 | 2,059.91 | 100.25 | 1.97 | 1,470.97 | 16.69 | 9,670.08 |
| 8 | 6,020.14 | 2,058.87 | 100.25 | 1.97 | 1,470.97 | 16.69 | 9,668.89 |
| 9 | 5,587.40 | 2,509.70 | 162.27 | 4.67 | 2,259.85 | 20.58 | 10,544.48 |
| 10 | 5,231.38 | 1,858.09 | 58.65 | 1.29 | 1,259.43 | 13.30 | 8,422.13 |
| 11 | 7,363.42 | 2,682.43 | 203.86 | 5.08 | 2,474.00 | 26.27 | 12,755.06 |
| 12 | 7,238.89 | 2,406.77 | 156.92 | 3.36 | 2,104.44 | 23.14 | 11,933.53 |

Based on the results of the simulations in Table 10, our first general observation is that in each scenario under consideration in which SOs are exchanged in the SM (ID = {5,6,7,8,10,12}), the acquired stocks are less in relation to the corresponding scenario in which SOs are not exchanged in the SM. (ID = {1,2,3,4,9,11}). More specifically, the Percentage Difference (PD) between corresponding scenarios is: PD (1,5) = 14.78%, PD (2,6) = 15.11%, PD (3,7) = 14.87%, PD (4,8) = 14.79%, PD (9,10) = 22.38% and PD (11,12) = 6.66%.

Moreover, it is essential to note that the most significant reduction in the acquisition of stocks at a penalty price, due to the activation of the SOs

exchange, is found between scenarios 9 and 10, in which the highest percentage of SOs (15%) is daily obtained. Respectively, the most negligible reduction is found between scenarios 11 and 12, at which the lowest percentage of SOs (5%) is obtained daily.

Table 11 - Cost of stocks acquired at penalty prices per scenario and OS (\$)

| Scn. ID | 1 | 2 | 3 | 4 | 5 | 6 | Total |
|---------|--------|-------|------|------|-------|------|--------|
| 1 | 208.72 | 69.93 | 5.49 | 0.20 | 67.83 | 0.84 | 353.01 |
| 2 | 207.36 | 69.28 | 5.35 | 0.20 | 66.59 | 0.82 | 349.59 |
| 3 | 206.64 | 68.91 | 5.24 | 0.20 | 66.35 | 0.80 | 348.15 |
| 4 | 206.52 | 68.89 | 5.24 | 0.19 | 66.21 | 0.79 | 347.85 |
| 5 | 201.41 | 58.40 | 3.02 | 0.09 | 43.29 | 0.64 | 306.86 |
| 6 | 201.25 | 55.87 | 3.01 | 0.08 | 42.40 | 0.61 | 303.22 |
| 7 | 200.98 | 55.64 | 3.01 | 0.08 | 42.37 | 0.61 | 302.69 |
| 8 | 200.98 | 55.61 | 3.01 | 0.08 | 42.37 | 0.61 | 302.65 |
| 9 | 185.20 | 68.07 | 4.87 | 0.20 | 65.42 | 0.75 | 324.50 |
| 10 | 173.46 | 49.87 | 1.76 | 0.05 | 36.02 | 0.48 | 261.64 |
| 11 | 246.68 | 72.84 | 6.12 | 0.21 | 71.78 | 0.96 | 398.59 |
| 12 | 242.60 | 65.16 | 4.71 | 0.14 | 61.13 | 0.85 | 374.58 |

Based on the results of the simulations in Table 11, in a similar analysis to that of Table 10, and regarding the reduction of the cost (due to the activation of the SOs exchange in the SM) that the OS pay for the acquisition of stocks at a penalty price, we have the following results: PD (1,5) = 13.99%, PD (2,6) = 14.21%, PD (3,7) = 13.97%, PD (4,8) = 13.89%, PD (9,10) = 21.45% and PD (11,12) = 6.21%. Again, the largest reduction in stock cost is found between scenarios 9 and 10, and the smallest reduction is between scenarios 11 and 12.

4.4.6 Reduction of the stocks acquisition at penalty prices due to the exchange of SOs in the SM.

As we have seen so far, the stocks bought at the BoD for each OS (Section 4.4.2), as well as the ones exchanged in the SM during each day (Section 4.4.3), remain unchanged as a quantity (GB) and as a cost (\$) in all 12 scenarios under consideration. The reason is that these functions occur daily, before the exercise of the SOs and before their exchange in the SM (exchange that takes place in six of the twelve scenarios we are considering).

However, the activation of the SOs exchange in the SM results in differences both in the number of stocks acquired through the exercise of the former (Section 4.4.4) and in the number of stocks that each OS is required to acquire at the EoD at a penalty price (Section 4.4.5).

This section compares these two metrics (established in sections 4.4.4 and 4.4.5) and

observes their behavior in total in each of the 12 scenarios under consideration: the six scenarios in which no SOs are exchanged with the corresponding six scenarios where SOs are exchanged between the OSs.

Table 12 - Comparison of stocks acquired through SOs and at penalty prices

| | Scn. 1 & 5 | Scn. 2 & 6 | Scn. 3 & 7 | Scn. 4 & 8 | Scn. 9 & 10 | Scn. 11 & 12 |
|-------------------|------------------|------------------|------------------|------------------|------------------|------------------|
| SM Off Penalty | 11,391.94 | 11,275.91 | 11,223.28 | 11,213.48 | 10,544.48 | 12,755.06 |
| SM Off via SOs | 3,488.80 | 3,604.83 | 3,657.46 | 3,667.26 | 4,336.26 | 2,125.68 |
| SM Off Sum | 14,880.74 | 14,880.74 | 14,880.74 | 14,880.74 | 14,880.74 | 14,880.74 |
| SM On Penalty | 9,823.91 | 9,691.54 | 9,670.08 | 9,668.89 | 8,422.13 | 11,933.53 |
| SM On via SOs | 5,056.83 | 5,189.20 | 5,210.66 | 5,211.85 | 6,458.61 | 2,947.21 |
| SM On Sum | 14,880.74 | 14,880.74 | 14,880.74 | 14,880.74 | 14,880.74 | 14,880.74 |

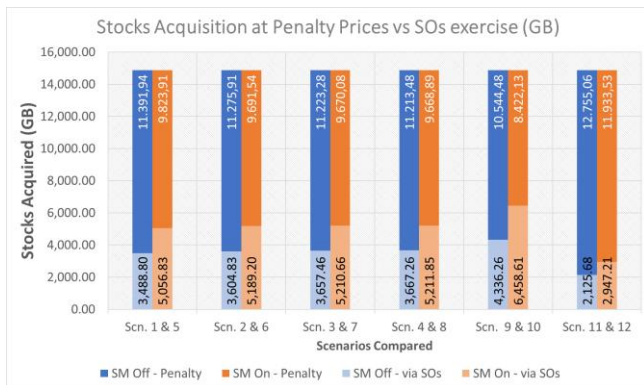


Fig. 7 - Stocks acquisition at penalty prices vs SOs exercise (GB)

As can be seen from Table 12 and Fig. 7, in each pair of correlated scenarios, activating the SOs exchange in the SM results in: (a) an increase in the stocks purchased through the exercise of SOs, and (b) a decrease in the number of stocks bought at penalty prices.

Regarding the costs that the OS pays, as depicted in Table 13 and Fig. 8, in each pair of related scenarios, the activation of the SOs exchange in the SM results in: (a) an increase in the costs of stocks purchased through the exercise of SOs, (b) a decrease in the cost of stocks bought at penalty prices, and (c) a small decrease in the cost of stocks acquired cumulatively with these two methods.

More specifically, the percentage differences of the reduction in the cumulative cost of stocks acquired by the OSs (and sold by the CDN), through SOs and at a penalty price, due to the activation of the

exchange of SOs in the SM, per pair of related scenarios are: PD (1,5) = 2.17%, PD (2,6) = 2.31%, PD (3,7) = 2.20%, PD (4,8) = 2.15%, PD (9,10) = 3.78%, PD (11,12) = 0.82%.

Table 13 - Comparison of costs of stocks acquired through SOs and at penalty prices

| | Scn. 1 & 5 | Scn. 2 & 6 | Scn. 3 & 7 | Scn. 4 & 8 | Scn. 9 & 10 | Scn. 11 & 12 |
|---------------------|---------------|---------------|---------------|---------------|---------------|---------------|
| SM Off Penalty | 353.01 | 349.59 | 348.15 | 347.85 | 324.50 | 398.59 |
| SM Off via SOs | 77.58 | 80.07 | 81.09 | 81.17 | 92.94 | 50.68 |
| SM Off Total | 430.58 | 429.66 | 429.24 | 429.02 | 417.44 | 449.27 |
| SM On Penalty | 306.86 | 303.22 | 302.69 | 302.65 | 261.64 | 374.58 |
| SM On via SOs | 114.48 | 116.62 | 117.20 | 117.23 | 140.31 | 71.04 |
| SM On Total | 421.34 | 419.83 | 419.89 | 419.88 | 401.95 | 445.62 |




Fig. 8 - Cost of stocks acquisition at penalty prices vs SOs exercise (\$)

Thus, the mechanism we introduce in this paper seems to have beneficial effects for OS for two reasons: (a) because it reduces their need for extra stocks acquisition (at a penalty price) from CDN, and (b) because it reduces the total cost they pay.

4.4.7 SOs and the results of their exchange in SM.

Let us review SOs and their functionality as a whole for each scenario under consideration. The OSs seek to preserve the right to purchase stocks in the future (at a lower than penalty price) by buying SOs. If they ultimately do not exercise some of the SOs, then the buying cost of unused SOs negatively affects their wallet. Also, their initial prediction was proven to be partially wrong. In the cases where SOs can be exchanged in the SM, OSs can sell excess SOs. So, the seller can recover part of their initial cost (of purchasing SOs), and the buyer can acquire stocks at standard (rather than penalty) prices, in addition to the cost of acquiring the SOs (which is relatively

low). The following results of this section are interesting to compare with the conclusions of Section 4.4.4.

We begin our review with a synopsis of the four basic SOs' metrics, per scenario, for all the OSs in total: Total (Purchased from CDN), Exercised, Not Exercised, and Exchanged (between OSs). The CDN is not involved in the SOs exchange process, as they are exchanged between the OSs. Also, any cost differences only concern OS's wallets (budgets).

Table 14 – Synopsis of the SOs used (GB) per scenario ID

| Scn. Id | Total | Exercised | Not Exercised | Exchanged |
|---------|----------|-----------|---------------|-----------|
| 1 | 4,259.68 | 3,488.80 | 770.88 | 0.00 |
| 2 | 3,754.77 | 3,604.83 | 149.94 | 0.00 |
| 3 | 3,711.85 | 3,657.46 | 54.39 | 0.00 |
| 4 | 3,673.29 | 3,667.26 | 6.03 | 0.00 |
| 5 | 5,530.84 | 5,056.83 | 474.00 | 1,547.45 |
| 6 | 5,238.58 | 5,189.20 | 49.38 | 1,549.60 |
| 7 | 5,211.85 | 5,210.66 | 1.19 | 1,656.81 |
| 8 | 5,211.85 | 5,211.85 | 0.00 | 1,656.81 |
| 9 | 6,003.13 | 4,336.26 | 1,666.86 | 0.00 |
| 10 | 7,572.55 | 6,458.61 | 1,113.94 | 2,097.24 |
| 11 | 2,188.22 | 2,125.68 | 62.53 | 0.00 |
| 12 | 2,961.40 | 2,947.21 | 14.19 | 866.59 |

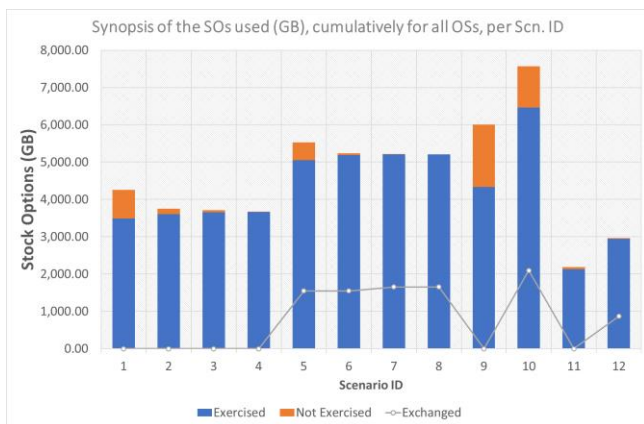


Fig. 9 – Synopsis of the SOs used (GB) per scenario ID

The results presented in Table 14 and Fig. 9 show that in each pair of related scenarios (with and without SOs exchange), activating the SOs exchange in the SM results in the following:

- An increase in the total number of SOs acquisitions (26.14% - 41.88%);
- An increase in the total number of SOs exercised and a decrease in the total number of SOs expired, both in absolute values and as a percentage of the total SOs acquired. In fact, in scenario 8, we observe a 100% practice of SOs.

We also notice that:

- **The duration of an SO** is inversely proportional to the probability that the SO will remain unused. Whether the SOs are exchanged at the SM or not, the longer they last, the less they expire. The duration of 6 or 7 days proved problematic because more rights are bought, but a much higher percentage of them expire compared to rights lasting 14+ days. Also, the increase in duration positively affects the percentage of those exchanged (2% - 4%).
- Increasing **the percentage of SOs bought at the BoD** results in an increase in the percentage of SOs that are exercised as well as the percentage of exchanged SOs. We need to consider a larger initial percentage, especially in combination with a duration longer than 6 days, a duration which, as we mentioned before, is problematic and leads to an increase in those that expire.

In total, we can say that, in terms of the specific characteristics of SOs, a purchase rate greater than or equal to 15% with a duration of 14 days or more is proposed.




Fig. 10 – Synopsis of the premium of SOs used (GB) per scenario ID

Regarding the premium (cost of acquiring) of the SOs, concerning the specific characteristics we discussed above (duration and percentage of acquisition) and how these affect the total costs and the cost of those SOs practiced and those that remained unused, the overall picture (Fig. 10) is almost identical.

4.4.8 Total stocks acquired and their respective cost

This section reviews the results of the 12 simulation scenarios in terms of the total stocks acquired and

their respective costs. It should be noted that the total cost paid by each OS (and collected by the CDN) also includes the cost of buying the SOs, regardless of whether the SOs were exercised.

Table 15 – Total stocks acquired per scenario (GB)

| Scn. Id | At the BoD & via SM | Via Exercising SOs | At the EoD | Total |
|---------|---------------------|--------------------|------------|------------|
| 1 | 138,293.36 | 3,488.80 | 11,391.94 | 153,174.10 |
| 2 | 138,293.36 | 3,604.83 | 11,275.91 | 153,174.10 |
| 3 | 138,293.36 | 3,657.46 | 11,223.28 | 153,174.10 |
| 4 | 138,293.36 | 3,667.26 | 11,213.48 | 153,174.10 |
| 5 | 138,293.36 | 5,056.83 | 9,823.91 | 153,174.10 |
| 6 | 138,293.36 | 5,189.20 | 9,691.54 | 153,174.10 |
| 7 | 138,293.36 | 5,210.66 | 9,670.08 | 153,174.10 |
| 8 | 138,293.36 | 5,211.85 | 9,668.89 | 153,174.10 |
| 9 | 138,293.36 | 4,336.26 | 10,544.48 | 153,174.10 |
| 10 | 138,293.36 | 6,458.61 | 8,422.13 | 153,174.10 |
| 11 | 138,293.36 | 2,125.68 | 12,755.06 | 153,174.10 |
| 12 | 138,293.36 | 2,947.21 | 11,933.53 | 153,174.10 |

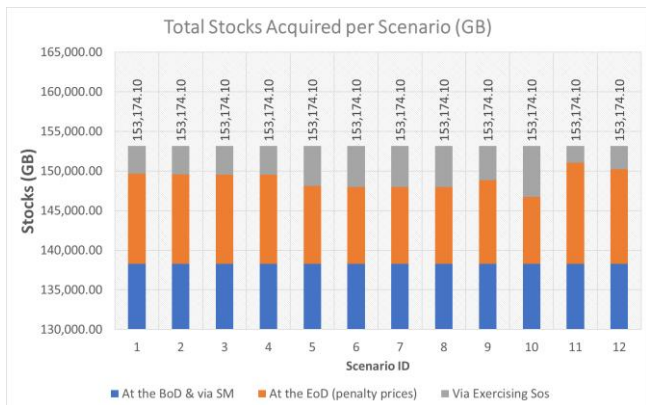


Fig. 11 – Total stocks acquired per scenario (GB)

According to what we have mentioned so far, and as evidenced by the results of Table 15 and Fig. 11, it is confirmed that the activation of the exchange of SOs in the SM, regardless of the scenario, reduces the stocks purchased at the EoD at penalty prices and partially “replaces” them with stocks acquired through exercising SOs at standard prices.

To take a closer look at these stocks’ “substitutions”, let us define PS (i) as the percentage of stocks acquired through SOs regarding the sum of the stocks acquired cumulatively through SOs and at a penalty price for the i_{th} scenario. Thus, in column “No SM”, we calculate the metric PS (i) for the scenarios in which the SOs are not exchanged in the SM, in column “SM” for the scenarios that are exchanged in the SM, while in column “PD” we calculate the percentage difference of PSs of the two compared scenarios.

Table 16 – Comparison of percentage of stocks acquired through SOs regarding the sum of those acquired through SOs and at penalty prices

| Scn. Id | PS (No SM) | Scn. Id | PS (SM) | PD |
|---------|------------|---------|---------|--------|
| 1 | 23.45% | 5 | 33.98% | 36.70% |
| 2 | 24.22% | 6 | 34.87% | 36.03% |
| 3 | 24.58% | 7 | 35.02% | 35.03% |
| 4 | 24.64% | 8 | 35.02% | 34.79% |
| 9 | 29.14% | 10 | 43.40% | 39.32% |
| 11 | 14.28% | 12 | 19.81% | 32.39% |

The best performance is observed in the scenarios with IDs 9 and 10. In these scenarios, (a) have the highest PS in both “No SM” and “SM” columns of Table 16, and (b) shows the most significant increase in the PS when activating the SM, which means that it shows the most considerable improvement in the percentage of replacement of “expensive” stocks by standard price stocks.

From the above, it is evident that the 15%, as a maximum percentage of initial acquisition of SOs per day, even if this pair of scenarios (9 and 10) has an SOs duration of just six days, is the main feature that gives this good result (as we have already mentioned in Section 4.4.4).

From there on, in every two comparative scenarios where the maximum percentage of initial acquisition of SOs is 10%, the performance is similar (Table 16 - column PD), with the better performance the one where the SOs duration is 7 days and a slightly declining performance as we approach the 28 days duration. The worst performance is presented in scenarios 11 and 12, where the maximum percentage of initial acquisition of SOs is 5%.

Table 17 - Total cost per scenario (\$)

| Scn. Id | At the BoD | Via Exercising SOs | At the EoD | SOs Premium | Total |
|---------|------------|--------------------|------------|-------------|----------|
| 1 | 3,828.80 | 77.58 | 353.01 | 8.21 | 4,267.60 |
| 2 | 3,828.80 | 80.07 | 349.59 | 7.31 | 4,265.77 |
| 3 | 3,828.80 | 81.09 | 348.15 | 7.27 | 4,265.31 |
| 4 | 3,828.80 | 81.17 | 347.85 | 7.26 | 4,265.08 |
| 5 | 3,828.80 | 114.48 | 306.86 | 10.07 | 4,260.22 |
| 6 | 3,828.80 | 116.62 | 303.22 | 9.64 | 4,258.28 |
| 7 | 3,828.80 | 117.20 | 302.69 | 9.70 | 4,258.39 |
| 8 | 3,828.80 | 117.23 | 302.65 | 9.81 | 4,258.50 |
| 9 | 3,828.80 | 92.94 | 324.50 | 10.71 | 4,256.96 |
| 10 | 3,828.80 | 140.31 | 261.64 | 12.64 | 4,243.39 |
| 11 | 3,828.80 | 50.68 | 398.59 | 4.77 | 4,282.85 |
| 12 | 3,828.80 | 71.04 | 374.58 | 6.28 | 4,280.70 |

The same results are also observed for the issue of cost in Table 17 and Fig. 12. In scenarios 9 and 10, there is a maximum reduction in the extra cost of buying stocks at penalty prices at the EoD. Part of this extra cost is replaced by lower (standard) stock purchase costs through the SOs exercise.



Fig. 12 - Total cost per scenario (\$)

As a common conclusion, we can say that the goal, which is none other than the maximum possible exemption from the purchase of expensive stocks (at the EoD at penalty prices), seems to be achieved by: (a) the further increase of the number of SOs purchased by each OS at the BoD and then exercised or exchanged in combination with (b) an SOs duration of 14 days or more.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we present an extension of the framework proposed and presented in our past work about the capital market metaphor in the area of CDNs. The previous version of the framework deals with the efficient prediction of resource consumption. **This framework version also utilizes an SM to exchange resources and the exercise of SOs to minimize the prediction misses that could not be remedied in the SM [1]–[3].**

The improvement proposed in this paper is the ability to exchange the available SOs between OSs through the existing SM.

Having developed a simulation program, we implement multiple year-long simulations of the proposed framework, observing the market's participants' conducted business. Based on the presented and analyzed results, it has been shown that **this ability has reduced the stocks that the OSs acquire at penalty prices and thus the cost they pay.** In addition, it has been shown that **a further increase, beyond the 15% percentage, in the number of SOs purchased from each OS at**

the BoD is needed, with a duration of 14 days or more.

In future work, we intend to design a mechanism that will regulate the two basic SOs parameters (number and duration) based on specific metrics such as the rate of reduction of the purchase of stocks at a penalty price, the number of SOs that expire unexploited and the total cost of acquired stocks. We also plan to simulate a market where the CDN has no extra resources to sell after those initially sold at the BoD. In this case, the SM will be the only way for the OSs to acquire extra stocks or SOs. The lack of the CDN's penalized price, serving until now as the maximum price of acquiring stocks, will probably affect the SM's selling prices based on the demand and competitiveness among the buyers and sellers.

DECLARATIONS

The authors declare that they have no conflict of interest. No funding was received to assist with the preparation of this manuscript.

AUTHOR CONTRIBUTIONS STATEMENT

- Elias Vathias: conceptualization, methodology, data curation, software, validation, writing, reviewing, and editing.
- Matthias Filippou Doukas: conceptualization, methodology, data curation, software, validation, and writing.
- Panagiotis Giannopoulos: conceptualization, methodology, data curation, software, validation, and writing.
- Stathes Hadjiefthymiades: conceptualization, methodology, writing, reviewing and editing, supervision.

REFERENCES

- [1] E. Vathias and S. Hadjiefthymiades, "A stock options metaphor for content delivery networks," *Netnomics*, Sep. 2022, doi: 10.1007/s11066-022-09153-7.
- [2] E. Vathias, D. Nikolopoulos, and S. Hadjiefthymiades, "A Capital Market Metaphor for Content Delivery Network Resources," in *2016 IEEE 30th International Conference on Advanced Information Networking and Applications (AINA)*, Mar. 2016, pp. 101–108. doi: 10.1109/AINA.2016.108.

- [3] E. Vathias, E. Katsarou, and S. Hadjiefthymiades, "A secondary market metaphor for content delivery networks," *Netnomics*, vol. 18, no. 2, Art. no. 2, Dec. 2017, doi: 10.1007/s11066-017-9120-z.
- [4] J. C. Hull, *Options, futures and other derivatives*. 2019.
- [5] S. A. Ross, *Corporate finance*, Twelfth edition. New York, NY: McGraw-Hill Education, 2019.
- [6] M. do R. Grossinho, Y. F. Kord, and D. Sevcovic, "Pricing American Call Options by the Black-Scholes Equation with a Nonlinear Volatility Function," Jul. 2017, Accessed: Jun. 06, 2021. [Online]. Available: <https://arxiv.org/abs/1707.00358v3>
- [7] M. Alghalith, "Pricing the American options using the Black-Scholes pricing formula," *Physica A: Statistical Mechanics and its Applications*, vol. 507, pp. 443-445, Oct. 2018, doi: 10.1016/j.physa.2018.05.087.
- [8] G. Barone-Adesi and R. E. Whaley, "Efficient Analytic Approximation of American Option Values," *The Journal of Finance*, vol. 42, no. 2, Art. no. 2, Jun. 1987, doi: 10.1111/j.1540-6261.1987.tb02569.x.
- [9] N. Ju and R. Zhong, "An Approximate Formula for Pricing American Options," *JOD*, vol. 7, no. 2, Art. no. 2, Nov. 1999, doi: 10.3905/jod.1999.319140.
- [10] K. Xu, Y. Zhong, and H. He, *Internet Resource Pricing Models*. New York, NY: Springer New York, 2014. doi: 10.1007/978-1-4614-8409-7.
- [11] J. K. MacKie-Mason, H. R. Varian, and others, "Pricing the internet," *Public access to the Internet*, vol. 269, p. 273, 1995.
- [12] L. W. McKnight and J. P. Bailey, Eds., *Internet economics*, 5. print. Cambridge, Mass.: MIT Press, 1999.
- [13] A. Odlyzko, "Internet pricing and the history of communications," *Computer Networks*, vol. 36, no. 5-6, Art. no. 5-6, Aug. 2001, doi: 10.1016/S1389-1286(01)00188-8.
- [14] M. Currence, A. Kurzon, D. Smud, and L. Trias, "A Causal Analysis of Usage-Based Billing on IP Networks," 2000. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.41.2035>
- [15] R. J. Edell, N. McKeown, and P. P. Varaiya, "Billing users and pricing for TCP," *IEEE J. Select. Areas Commun.*, vol. 13, no. 7, Art. no. 7, Sep. 1995, doi: 10.1109/49.414636.
- [16] M. L. Honig and K. Steiglitz, "Usage-based pricing of packet data generated by a heterogeneous user population," in *Proceedings of INFOCOM'95*, Boston, MA, USA, 1995, pp. 867-874. doi: 10.1109/INFCOM.1995.515958.
- [17] R. Edell and P. Varaiya, "Providing Internet access: what we learn from INDEX," *IEEE Network*, vol. 13, no. 5, Art. no. 5, Oct. 1999, doi: 10.1109/65.793687.
- [18] J. Crowcroft and P. Oechslin, "Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP," *SIGCOMM Comput. Commun. Rev.*, vol. 28, no. 3, Art. no. 3, Jul. 1998, doi: 10.1145/293927.293930.
- [19] R. J. Gibbens and F. P. Kelly, "Resource pricing and the evolution of congestion control," *Automatica*, vol. 35, no. 12, Art. no. 12, Dec. 1999, doi: 10.1016/S0005-1098(99)00135-1.
- [20] F. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecomm.*, vol. 8, no. 1, Art. no. 1, Jan. 1997, doi: 10.1002/ett.4460080106.
- [21] N. Keon and G. "Anand" Anandalingam, "A New Pricing Model for Competitive Telecommunications Services Using Congestion Discounts," *INFORMS Journal on Computing*, vol. 17, no. 2, Art. no. 2, May 2005, doi: 10.1287/ijoc.1030.0062.
- [22] S. Kunniyur and R. Srikant, "End-to-end congestion control schemes: utility functions, random losses and ecn marks," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, Art. no. 5, Oct. 2003, doi: 10.1109/TNET.2003.818183.
- [23] J. K. MacKie-Mason, "A Smart Market for Resource Reservation in a Multiple Quality of Service Information Network," *SSRN Journal*, 1997, doi: 10.2139/ssrn.975871.
- [24] J. K. MacKie-Mason and H. R. Varian, "Pricing congestible network resources," *IEEE J. Select. Areas Commun.*, vol. 13, no. 7, Art. no. 7, Sep. 1995, doi: 10.1109/49.414634.

- [25] M. Yuksel and S. Kalyanaraman, "Pricing granularity for congestion-sensitive pricing," in *Proceedings of the Eighth IEEE Symposium on Computers and Communications. ISCC 2003*, Kemer-Antalya, Turkey, 2003, pp. 169–174. doi: 10.1109/ISCC.2003.1214118.
- [26] R. Cocchi, D. Estrin, S. Shenker, and L. Zhang, "A study of priority pricing in multiple service class networks," in *Proceedings of the conference on Communications architecture & protocols - SIGCOMM '91*, Zurich, Switzerland, 1991, pp. 123–130. doi: 10.1145/115992.116005.
- [27] R. Cocchi, S. Shenker, D. Estrin, and Lixia Zhang, "Pricing in computer networks: motivation, formulation, and example," *IEEE/ACM Trans. Networking*, vol. 1, no. 6, Art. no. 6, Dec. 1993, doi: 10.1109/90.266050.
- [28] A. Odlyzko, "Paris metro pricing for the internet," in *Proceedings of the 1st ACM conference on Electronic commerce - EC '99*, Denver, Colorado, United States, 1999, pp. 140–147. doi: 10.1145/336992.337030.
- [29] M. H. Dahshan and P. K. Verma, "Resource Based Pricing Framework for Integrated Services Networks," *JNW*, vol. 2, no. 3, Art. no. 3, Jun. 2007, doi: 10.4304/jnw.2.3.36-45.
- [30] C. Dovrolis, D. Stiliadis, and P. Ramanathan, "Proportional differentiated services: delay differentiation and packet scheduling," *IEEE/ACM Trans. Networking*, vol. 10, no. 1, Art. no. 1, Feb. 2002, doi: 10.1109/90.986503.
- [31] G. Fankhauser and B. Plattner, *Diffserv Bandwidth Brokers as Mini-Markets*. 1999.
- [32] A. Gupta, D. O. Stahl, and A. B. Whinston, "An economic approach to networked computing with priority classes," *Journal of Organizational Computing and Electronic Commerce*, vol. 6, no. 1, Art. no. 1, Jan. 1996, doi: 10.1080/10919399609540269.
- [33] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar, "Pricing, provisioning and peering: dynamic markets for differentiated Internet services and implications for network interconnections," *IEEE J. Select. Areas Commun.*, vol. 18, no. 12, Art. no. 12, Dec. 2000, doi: 10.1109/49.898733.
- [34] Xin Wang and H. Schulzrinne, "Pricing network resources for adaptive applications in a differentiated services network," in *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No.01CH37213)*, Anchorage, AK, USA, 2001, vol. 2, pp. 943–952. doi: 10.1109/INFCOM.2001.916286.
- [35] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, no. 3, Art. no. 3, Mar. 1998, doi: 10.1057/palgrave.jors.2600523.
- [36] Mung Chiang, Shengyu Zhang, and P. Hande, "Distributed rate allocation for inelastic flows: optimization frameworks, optimality conditions, and optimal algorithms," in *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, Miami, FL, USA, 2005, vol. 4, pp. 2679–2690. doi: 10.1109/INFCOM.2005.1498551.
- [37] T. Başar and R. Srikant, "A Stackelberg Network Game with a Large Number of Followers," *Journal of Optimization Theory and Applications*, vol. 115, no. 3, Art. no. 3, Dec. 2002, doi: 10.1023/A:1021294828483.
- [38] T. Basar and R. Srikant, "Revenue-maximizing pricing and capacity expansion in a many-users regime," in *Proceedings. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, New York, NY, USA, 2002, vol. 1, pp. 294–301. doi: 10.1109/INFCOM.2002.1019271.
- [39] R. B. Myerson, *Game theory: analysis of conflict*, 6. print. Cambridge, Mass.: Harvard Univ. Press, 2004.
- [40] J. Nash, "Non-Cooperative Games," *The Annals of Mathematics*, vol. 54, no. 2, Art. no. 2, Sep. 1951, doi: 10.2307/1969529.
- [41] G. Alyfantis, S. Hadjiefthymiades, L. Merakos, and P. Kostopoulos, "A distributed algorithm for sharing Web cache disk capacity," in *12th International Conference on Parallel and Distributed Systems - (ICPADS'06)*, Minneapolis, MN, USA, 2006, p. 8 pp. doi: 10.1109/ICPADS.2006.12.


- [42] K. Hosanagar, R. Krishnan, J. Chuang, and V. Choudhary, "Pricing and Resource Allocation in Caching Services with Multiple Levels of Quality of Service," *Management Science*, vol. 51, no. 12, Art. no. 12, Dec. 2005, doi: 10.1287/mnsc.1050.0420.
- [43] K. Hosanagar, R. Krishnan, J. Chuang, and V. Choudhary, "Optimal Pricing and Capacity Allocation in Vertically Differentiated Web Caching Services," *ICIS 2002 Proceedings*, Dec. 2002, [Online]. Available: <https://aisel.aisnet.org/icis2002/44>
- [44] N. Laoutaris, O. Telelis, V. Zissimopoulos, and I. Stavrakakis, "Distributed Selfish Replication," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 12, Art. no. 12, Dec. 2006, doi: 10.1109/TPDS.2006.171.
- [45] N. Laoutaris, G. Smaragdakis, A. Bestavros, and I. Stavrakakis, "Mistreatment in Distributed Caching Groups: Causes and Implications," in *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, Barcelona, Spain, 2006, pp. 1–13. doi: 10.1109/INFOCOM.2006.131.
- [46] O. Ercetin and L. Tassiulas, "Pricing strategies for differentiated services content delivery networks," *Computer Networks*, vol. 49, no. 6, Art. no. 6, Dec. 2005, doi: 10.1016/j.comnet.2005.03.001.
- [47] K. Hosanagar, R. Krishnan, M. Smith, and J. Chuang, "Optimal pricing of content delivery network (CDN) services," in *37th Annual Hawaii International Conference on System Sciences, 2004. Proceedings of the*, Big Island, HI, USA, 2004, p. 10 pp. doi: 10.1109/HICSS.2004.1265480.
- [48] K. Hosanagar, J. Chuang, R. Krishnan, and M. D. Smith, "Service Adoption and Pricing of Content Delivery Network (CDN) Services," *Management Science*, vol. 54, no. 9, Art. no. 9, Sep. 2008, doi: 10.1287/mnsc.1080.0875.
- [49] O. Ercetin and L. Tassiulas, "Market-based resource allocation for content delivery in the internet," *IEEE Trans. Comput.*, vol. 52, no. 12, Art. no. 12, Dec. 2003, doi: 10.1109/TC.2003.1252853.
- [50] J. Broberg, R. Buyya, and Z. Tari, "MetaCDN: Harnessing 'Storage Clouds' for high performance content delivery," *Journal of Network and Computer Applications*, vol. 32, no. 5, Art. no. 5, Sep. 2009, doi: 10.1016/j.jnca.2009.03.004.
- [51] M. Pathan, J. Broberg, and R. Buyya, "Maximizing Utility for Content Delivery Clouds," in *Web Information Systems Engineering - WISE 2009*, vol. 5802, G. Vossen, D. D. E. Long, and J. X. Yu, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 13–28. Accessed: Mar. 28, 2021. [Online]. Available: http://link.springer.com/10.1007/978-3-642-04409-0_7
- [52] V. K. Adhikari *et al.*, "Unreeling netflix: Understanding and improving multi-CDN movie delivery," in *2012 Proceedings IEEE INFOCOM*, Orlando, FL, USA, Mar. 2012, pp. 1620–1628. doi: 10.1109/INFOCOM.2012.6195531.
- [53] M. K. Mukerjee, I. N. Bozkurt, B. Maggs, S. Seshan, and H. Zhang, "The Impact of Brokers on the Future of Content Delivery," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, Atlanta GA USA, Nov. 2016, pp. 127–133. doi: 10.1145/3005745.3005749.
- [54] M. K. Mukerjee, I. N. Bozkurt, D. Ray, B. M. Maggs, S. Seshan, and H. Zhang, "Redesigning CDN-Broker Interactions for Improved Content Delivery," in *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, Incheon Republic of Korea, Nov. 2017, pp. 68–80. doi: 10.1145/3143361.3143366.
- [55] C. Canali, A. Corbelli, and R. Lancellotti, "Designing a Private CDN with an Off-Sourced Network Infrastructure: Model and Case Study," in *2018 26th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Split, Sep. 2018, pp. 1–6. doi: 10.23919/SOFTCOM.2018.8555766.
- [56] H. Lee, D. Lee, and Y. Yi, "On the economic impact of Telco CDNs and their alliance on the CDN market," in *2014 IEEE International Conference on Communications (ICC)*, Sydney, NSW, Jun. 2014, pp. 2950–2955. doi: 10.1109/ICC.2014.6883773.

- [57] H. Lee, L. Duan, and Y. Yi, "On the competition of CDN companies: Impact of new telco-CDNs' federation," in *2016 14th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, Tempe, AZ, May 2016, pp. 1–8. doi: 10.1109/WIOPT.2016.7492919.
- [58] P. A. Frangoudis, L. Yala, A. Ksentini, and T. Taleb, "An architecture for on-demand service deployment over a telco CDN," in *2016 IEEE International Conference on Communications (ICC)*, Kuala Lumpur, Malaysia, May 2016, pp. 1–6. doi: 10.1109/ICC.2016.7510921.
- [59] S. Spagna *et al.*, "Design principles of an operator-owned highly distributed content delivery network," *IEEE Commun. Mag.*, vol. 51, no. 4, Art. no. 4, Apr. 2013, doi: 10.1109/MCOM.2013.6495772.
- [60] P. A. Frangoudis, L. Yala, and A. Ksentini, "CDN-As-a-Service Provision Over a Telecom Operator's Cloud," *IEEE Trans. Netw. Serv. Manage.*, vol. 14, no. 3, Art. no. 3, Sep. 2017, doi: 10.1109/TNSM.2017.2710300.
- [61] D. Xu, S. S. Kulkarni, C. Rosenberg, and H.-K. Chai, "Analysis of a CDN-P2P hybrid architecture for cost-effective streaming media distribution," *Multimedia Systems*, vol. 11, no. 4, Art. no. 4, Apr. 2006, doi: 10.1007/s00530-006-0015-3.
- [62] H. Yin *et al.*, "Design and deployment of a hybrid CDN-P2P system for live video streaming: experiences with LiveSky," in *Proceedings of the seventeen ACM international conference on Multimedia - MM '09*, Beijing, China, 2009, p. 25. doi: 10.1145/1631272.1631279.
- [63] I. A. Kash, R. Murty, and D. C. Parkes, "Enabling Spectrum Sharing in Secondary Market Auctions," *IEEE Trans. on Mobile Comput.*, vol. 13, no. 3, Art. no. 3, Mar. 2014, doi: 10.1109/TMC.2013.17.
- [64] M. Bitsaki, G. D. Stamoulis, and C. Courcoubetis, "A new strategy for bidding in the network-wide progressive second price auction for bandwidth," in *Proceedings of the 2005 ACM conference on Emerging network experiment and technology - CoNEXT'05*, Toulouse, France, 2005, p. 146. doi: 10.1145/1095921.1095941.
- [65] M. Bitsaki, G. D. Stamoulis, and C. Courcoubetis, "An efficient auction-based mechanism for hierarchically structured bandwidth markets," *Computer Communications*, vol. 29, no. 7, Art. no. 7, Apr. 2006, doi: 10.1016/j.comcom.2005.08.012.

AUTHORS



Elias Vathias received his B.Sc. and M.Sc. in informatics and telecommunications from the Department of Informatics & Telecommunications (DIT) of the National and Kapodistrian University of Athens (NKUA), Athens, Greece. Currently, he is a Ph.D. candidate in the same department and works as an IT consultant and database administrator in the public sector of Greece. His research interests are in the areas of information systems (business informatics), distributed computing, and computer communications (networks). His most recent publication is 'A Stock Options Metaphor for Content Delivery Networks'.



Matthias Filippou Doukas received his B.Sc. Degree in informatics and telecommunications from the National and Kapodistrian University of Athens (NKUA) in 2015 and a master's degree in management and economics of telecommunication networks from the same university in 2019. He has been working as an IT systems engineer at Alpha Bank since 2018. Previously, from 2015 to 2018, he worked as an IT systems engineer at AMPLUS.



Panagiotis Giannopoulos received his B.Sc. degree in informatics and telecommunications from the National and Kapodistrian University of Athens (NKUA) in 2014 and a master's degree in management and economics of telecommunication networks from the same university in 2019. Since 2014, he works as a consultant/software developer at Accenture.



Stathes Hadjiethymiades received his B.Sc., M.Sc., and Ph.D. in informatics and telecommunications from the Department of Informatics & Telecommunications (DIT) of the University of Athens (UoA), Athens, Greece. He also received a joint engineering-economics M.Sc. degree

from the National Technical University of Athens. In 1992 he joined the Greek consulting firm Advanced Services Group, Ltd., as an analyst/developer of telematic applications and systems. In 1995 he became a member of the Communication Networks Laboratory of UoA. During the period 2001–2002, he served as a visiting assistant professor at the University of Aegean, Department of Information and Communication Systems Engineering. In 2002 he joined the faculty of the Hellenic Open University (Department of Informatics), Patras, Greece, as an assistant professor. Since the beginning of 2004, he belongs to the faculty of UoA, DIT, in the Sector of Computer Systems and Applications, where he is presently a professor. He has participated in numerous projects realized in the context of EU programs and national initiatives. His research interests are in the areas of distributed, mobile and pervasive computing, web systems engineering, Internet of things, and networked multimedia applications and systems. He is the author of over 300 publications in these areas.

APPENDIX

A1 Model formulation

The prediction scheme is realized by means of the function $Q(w, t)$, where w denotes a history window measured in days that is taken into account for the prediction (in our case $w=7$, a full week), t is the prediction horizon (i.e. the predicted resources will be needed in time $c+t$, where c stands from the current time).

A resource reservation is a quantity x which is detached from the pool of resources of the CDN for a period of T days (1 day).

Then the range of the Q function is N (integer). $Q(w, t) \in N$. The volume $Q(w, t)$ of resources reserved by the OS is expressed in bytes (i.e. we adopt a fine-grained resource management).

We also monitor and manage the surplus (or deficit) of the reserved resources that can be further traded through the SM. To quantify the surplus (or deficit), we establish a metric indicating the actual use of resources throughout the day. This is denoted by $L(t)$, where t is the time of prediction deficiency quantification.

$$Q(w, t - \Delta t) - L(t) = S(t)$$

Should $S(t)$ be positive, a surplus is manifested, and the unused resources will be made available through a secondary Stock Market (SM). The volume-wise size of the market at time t is represented by $SM(t)$.

$$SM(t) = \sum_{\text{servers } i}^{all\ origin} (Q_i(w, t - \Delta t) - L_i(t)). a = \sum_{\text{servers } i}^{all\ origin} S_i(t)$$

A percentage (α) of the surplus is committed to the SM. This percentage is in the order of 100%.

Similarly, to the stocks case, an OS may opt for purchasing SOs that would allow the mitigation of prediction failures. An SO refers to the future purchase and use of y resources from the CDN at a price that is a priori defined. This pricing takes place in time t for a maturity period m . The Black Scholes call option formula is usually used for the pricing and is calculated by multiplying the stock price by the cumulative standard normal probability distribution function. Thereafter, the Net Present Value (NPV) of the strike price multiplied by the cumulative standard normal distribution is subtracted from the previous calculation's resulting value.

In mathematical notation:

$$C(m) = S_t N(d_1) - K e^{-rm} N(d_2)$$

where

$$d_1 = \frac{\ln \frac{S_t}{K} + \left(r + \frac{\sigma^2}{2} \right) m}{\sigma \sqrt{m}}$$

and

$$d_2 = d_1 - \sigma \sqrt{m}$$

where:

C = Call option price (output)

S = Current stock (or other underlying) price

K = Strike or exercise price (in our case is the price that CDN sells – Table 3)

r = Risk-free interest rate (0 in our case)

m = Time to maturity (expressed in years)

N = The cumulative distribution function of the standard normal distribution

σ = The volatility of returns of the underlying asset

Similarly to the demand prediction – resource reservation chain, the $Q(w, t)$ formula output is also exploited for scaling the SOs claims that the OS addresses the CDN. $Q(w, t)$ is scaled by a factor b ($0 < b < 1$), which indicates the number of SOs that the OS tries to secure in order to take corrective measures against prediction failures. In this paper $b=0.15$ (15%). As explained in Section 4.1, the type of SOs that we have adopted is European, but the respective rights can be exercised at any time until their maturity milestone since no dividends are paid in our case.

Similarly, to the existing resources, SOs can exhibit a surplus (SOS) when acquired by a specific OS. $Y(t)$ indicates this SOS (volume-wise) at the end of a given day, t . SOS can be viewed as a set of doublets:

$$SOS(t) = \{(O_i, m_i)_i\}$$

where O_i denotes the volume of SOs owned, acquired by the OS in the past not exercised until the present time t and m_i the respective maturity period. The OS trades through the SM the SOS members with the nearest maturity dates (min m_i) [shortest maturity first].

At any given time v , the Resources Reserved (RR) by a specific OS (Ok) are the cumulation of reservations and discharging (or acquisition) of resources through the SM (surplus or deficit) for the entire period of system operation.

$$RR_k(v) = \sum_{t=0}^v [Q_k(\cdot, t) + S_k(t)]$$