

Stochastic Planning for Content Delivery: Unveiling the Benefits of Network Functions Virtualization

Michele Mangili^{*†}, Fabio Martignon^{*‡} and Antonio Capone[†]

^{*} LRI, Université Paris-Sud
{michele.mangili, fabio.martignon}@lri.fr

[†] DEIB, Politecnico di Milano
capone@elet.polimi.it

[‡] IUF, Institut Universitaire de France

Abstract—Content Delivery Networks (CDNs) have been identified as one of the relevant use cases where the emerging paradigm of Network Functions Virtualization (NFV) will likely be beneficial. In fact, virtualization fosters flexibility, since on-demand resource allocation of virtual CDN nodes can accommodate sudden traffic demand changes. However, there are cases where physical appliances should still be preferred, therefore we envision a mixed architecture in between these two solutions, capable to exploit the advantages of both of them. Motivated by these reasons, in this paper we formulate a two-stage stochastic planning model that can be used by CDN operators to compute the optimal long-term network planning decision, deploying physical CDN appliances in the network and/or leasing resources for virtual CDN nodes in data centers. Key findings demonstrate that for a large range of pricing options and traffic profiles, NFV can significantly save network costs spent by the operator to provide the content distribution service.

Keywords—Content Distribution, Network Function Virtualization, Resource Allocation, Stochastic Network Planning.

I. INTRODUCTION

The worldwide success of content-rich web applications like social networks or on-demand streaming services has forced network operators to invest a significant amount of money in order to keep up-to-date their communication infrastructures [1]. In particular, Content Delivery Networks (CDNs) have nowadays become a necessary (and well-established) technology to efficiently serve the traffic demands that consumers are generating, while supporting the high level of performance and reliability that providers are demanding [2].

Although CDN is an effective infrastructure to move replicas of popular contents closer to the users' locations, it requires significant investments to be built and operated. As an example, the Akamai infrastructure comprises more than 61 000 servers deployed in 1 000 networks and 70 countries worldwide [3]. To reduce the capital expenditures and improve the performance of CDNs, organizations such as the Internet Engineering Task Force (IETF) and the European Telecommunications Standards Institute (ETSI) have recently begun a standardization process for two alternative architectures:

- Content Delivery Network Interconnection (CDNI);
- Virtual Content Delivery Network (vCDN).

Despite the fact that they both have to deal with network content distribution, these proposals have a radically different scope: the former (CDNI) is mostly concerned with the co-operation of many CDN providers [4], whereas the latter (vCDN) proposes to virtualize the CDN services on

top of the novel layer for Network Functions Virtualization (NFV) [5].

While both the architectures aim at optimally exploiting available physical resources, the grounds of CDNI are settled on agreements between different CDN operators that often are in direct competition in the same market. On the other hand, the NFV approach is to run network functions in a virtualized environment, executed on a shared physical infrastructure composed of industry standard high volume servers, storage and switches [5]. Therefore, vCDN implemented on top of NFV enjoys the positive advantage of avoiding potential competition issues, since the virtualized environment ensures the necessary level of isolation between the different network functions. Furthermore, spare NFV substrate capacity can be leased by network operators to third parties, a condition that makes vCDN appear even more profitable.

Motivated by the previous background, in this paper we tackle a fundamental issue that arises in such context: the planning problem for a mixed physical-virtual Content Delivery Network under uncertain traffic demands. In our formulation, the CDN operator can choose between *purchasing* physical CDN appliances and *leasing* instances of virtual CDN nodes provided by an infrastructure operator. However, while vCDN nodes can be activated on-demand if the traffic requests require to do so, the installation of physical CDN nodes must be chosen on a long-term schedule. For both physical or virtual CDN surrogate servers, the operator must carefully choose their location, while minimizing the overall costs. Due to the fact that planning is performed on a long-term basis, the theoretical framework of *stochastic optimization* will be used to guarantee robustness of the solution with respect to the uncertainty in the probabilistic description of future traffic demands.

The contribution of this paper is summarized as follows:

- 1) We formulate a two-stage stochastic planning model used by CDN operators to compute the optimal, long-term network planning decision, under traffic demands uncertainty.
- 2) We propose a greedy heuristic approach that finds good solutions (close to the optimum, in several cases) even for large-scale network topologies, and we compare its execution time with *exact* solution strategies: (1) the deterministic equivalent program in the extensive form and (2) the L-shaped algorithm (*single* and *multicut* versions) [6].
- 3) We perform an extensive numerical evaluation, considering real scale topologies and a wide range of parameters.

Our key findings suggest that a mixed physical-virtual CDN infrastructure leads to significant lower costs when compared to those obtained by a standard CDN, while being robust with respect to sudden traffic demand changes.

The paper is organized as follows: Sec. II discusses related work. In Sec. III we present our contribution; in particular, the system is described in Sec. III-A; the optimization model is presented in Sec. III-B, while the greedy algorithm is illustrated in Sec. III-C. Numerical results are presented in Sec. IV, and Sec. V concludes this paper.

II. RELATED WORK

In this section, we survey relevant literature on Network Functions Virtualization (Sec. II-A), Content Delivery Networks (Sec. II-B) and Stochastic Optimization techniques (Sec. II-C).

A. Network Functions Virtualization

To support the network virtualization paradigm, one of the challenges that must be solved is to find a mapping between a set of requests for virtual network resources and the available underlying physical infrastructure, ensuring that desired performance requirements on nodes and links are guaranteed [7]. This is the *virtual network embedding* (VNE) problem, which is known to be NP-hard, since it can be reduced to the multi-way separator problem [8]. VNE has received a lot of attention from the community, and several heuristic algorithms have been proposed, e.g. in [9]–[11].

Jarray et al. propose in [9] a column-generation technique coupled with a rounding heuristic to discover the most profitable embedding, under the constrained physical capacity of the infrastructure. Deterministic and randomized rounding techniques are used by Chowdhury et al. in [10], where they further facilitate the virtual link mapping by designing an augmented graph description to efficiently support node location constraints. Cheng et al. in [11] solve the node mapping step with a greedy algorithm: higher ranking is given to the nodes that possess more spare resources and are placed in better locations of the network.

Rather than assuming that the operator knows *a-priori* the traffic demands, our contribution is to consider the case where their probability distribution is known, and our proposed formulation is robust with respect to such uncertainty.

B. Content Delivery Networks

In the last few years, *content multihoming* is emerging as a novel technique for content delivery networks that makes it possible to jointly use many CDN services: [1], [2], [12].

Adhikari et al. show in [1] that the Netflix infrastructure already leverages multiple CDNs (Akamai, LimeLight and Level-3). The authors observe that the customers are mapped to a particular CDN in a rather static manner. In [2], Liu et al. further confirm that other major content publishers such as Hulu, Microsoft, Apple, Facebook and MSNBC are currently already exploiting content multihoming. Furthermore, given the practical relevance of this architecture, they design optimization algorithms to minimize the overall distribution costs under constrained quality requirements. Finally, Wang

et al. extended in [12] the work of Liu et al., by explicitly considering capacity constraints on the surrogate nodes.

The novelty of our approach is to consider a new distribution architecture, implemented on top of NFV, where virtual CDN nodes can be used for content delivery purposes. Rather than forcing the provider to settle agreements with other CDN operators, our proposal guarantees better isolation and dramatically limits potential competition issues.

C. Stochastic Optimization

In practical scenarios, network design cannot assume that future traffic demands are known a-priori; on the other hand, more advanced optimization techniques must be used to take into consideration the stochastic nature of input parameters: [13]–[15].

Atamtürk et al. formulate in [13] a two-stage network design model with traffic demand uncertainty. In their approach, the operator performs the planning decision according to a probabilistic description of traffic demands. The value of the second-stage recourse variables is chosen by changing flow routing. Liu describes in [14] the basic stochastic procedures applied to a flow assignment network design problem, showing the *here-and-know* solution and the *scenario-tracking* result obtained for the flow-assignment. A multistage stochastic programming model for mobile radio access networks has been proposed by Eisenblätter et al. in [15]. In their formulation they jointly take into consideration the coverage and capacity of their communication infrastructure.

In line with previous literature, our formulation takes into account the uncertainty embedded in future traffic demands. Our contribution is to apply the theoretical framework of stochastic optimization to content distribution.

III. OPTIMAL CONTENT DELIVERY IN NFV

In this section we describe our proposed solution for the optimal content delivery planning in NFV. Sec. III-A introduces the system model and relevant assumptions. In Sec. III-B we formulate the optimization model, while in Sec. III-C we discuss the design of a near-optimal heuristic algorithm.

A. System Model and Assumptions

Figure 1 shows the system model we consider in our proposal. In this work we tackle the long-term planning problem from the point of view of a CDN provider. The aim of the provider is to perform two choices:

- 1) Select *whether* and *where* physical CDN nodes should be installed in the network topology;
- 2) Select the optimal *request routing*, given the installed physical CDNs and the virtual nodes available.

Since the planning decision is operated on a long-term time schedule, the provider does not deterministically know what is going to happen in the future. On the other hand, we assume that an estimate of the continuous probability distribution of future traffic demands is known for the planning problem. However, for the sake of simplicity, and as frequently done in the literature (e.g. [6], [15]), we discretize this information

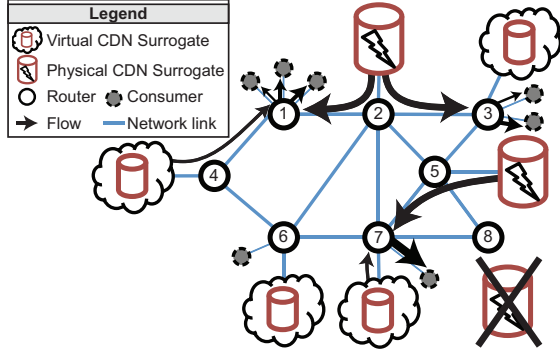


Figure 1. *System model*. The network is composed by consumers, routers, virtual and physical CDN surrogates. Our proposed optimization model selects (1) the planning of physical CDNs and (2) request routing.

on a finite number of scenarios, therefore our traffic model is jointly *time-varying* and *stochastic* in its nature.

As shown in Fig. 1, traffic demands are expressed by the consumers. Virtual and physical CDN nodes can both be used to efficiently serve consumers' demands, however there exist major differences (in terms of capacity, activation choice and pricing policy) between these two types of CDN hosts:

- *Capacity*: virtual CDN nodes can serve a lower amount of traffic requests since the presence of the hypervisor and the shared hardware infrastructure reduces the throughput of the CDN surrogates.
- *Activation choice*: virtual CDN servers can be used on-demand, and they do not need to be explicitly activated. On the other hand, if the operator chooses to install a physical CDN server, it will be activated once and it will stay active throughout the entire time horizon.
- *Pricing policy*: physical CDN nodes have an activation price that considers both the capital expenditure (CAPEX) for the acquisition of the device as well as the long-term operational expenditure (OPEX) costs. On the other hand, the virtual CDN nodes have a traffic-proportional price related to the OPEX cost component, which is the per-bandwidth leasing price that the vCDN owner charges.

Since virtual CDN nodes can be used on-demand, they can serve the portion of traffic requests with the highest variability. To improve the quality of service of a CDN, surrogate servers must be selected close to the location of consumers. For this purpose, we use the link delay to control the performance of the infrastructure: we assume that the content provider wants to serve a fraction of the overall requests within a bounded limit on the delay.

Popular CDN providers such as Amazon CloudFront or Microsoft Azure CDN do not have an activation cost but charge for their services according to the amount of traffic that surrogates are providing, regardless of the caching storage used. Moreover, frequent flash crowds make the object popularity suddenly change, whereas having an estimate of the aggregate future demands is instead much easier [16]. For these reasons, we focus on infrastructure planning and request routing, while we do not tackle the replica placement problem.

Table I. NOTATION USED IN THIS PAPER.

Input Parameters	
\mathcal{D}	Set of consumers (<i>destination nodes</i>)
\mathcal{S}	Set of candidate surrogate servers (<i>source nodes</i>) $\mathcal{S} = \mathcal{S}^P \cup \mathcal{S}^V$
\mathcal{S}^P	Set of candidate physical CDN servers
\mathcal{S}^V	Set of candidate virtual CDN servers
\mathcal{T}	Set of time slots
Φ	Set of stochastic scenarios
$r_d^{t,\phi}$	Traffic requests of client $d \in \mathcal{D}$, at time slot $t \in \mathcal{T}$, for scenario $\phi \in \Phi$
ϵ	Minimum service level guaranteed (fraction of traffic requests served with a bounded delay of at most Δ)
Δ	Maximum tolerated delay
$\delta_{s,d}$	Delay between the nodes $s \in \mathcal{S}$, $d \in \mathcal{D}$
\mathcal{K}_s^P	Capacity of the physical CDN server $s \in \mathcal{S}^P$
\mathcal{K}_s^V	Capacity of the virtual CDN server $s \in \mathcal{S}^V$
\mathcal{C}_s^P	CAPEX and OPEX costs of the physical CDN server installed at a candidate site $s \in \mathcal{S}^P$
\mathcal{C}_s^V	Usage cost of the virtual CDN server at the candidate site $s \in \mathcal{S}^V$
p_ϕ	Realization probability for the scenario $\phi \in \Phi$

Decision Variables	
a_s	0-1 Physical CDN activation variable. $a_s = 1$ if a physical CDN is installed at the candidate point $s \in \mathcal{S}^P$
$y_{s,d}^{t,\phi}$	Physical CDN flow variable for requests served by $s \in \mathcal{S}^P$ to client $d \in \mathcal{D}$, at time $t \in \mathcal{T}$, and scenario $\phi \in \Phi$
$z_{s,d}^{t,\phi}$	Virtual CDN flow variable for requests served by $s \in \mathcal{S}^V$ to client $d \in \mathcal{D}$, at time $t \in \mathcal{T}$, and scenario $\phi \in \Phi$

B. Optimization Model

In this section we describe the optimization model we formulate for the optimal planning of a mixed physical-virtual CDN infrastructure. The notation is summarized in Table I.

Let $\mathcal{S} = \mathcal{S}^P \cup \mathcal{S}^V$ be the set of CDN surrogate servers, where \mathcal{S}^P and \mathcal{S}^V represent candidate physical and virtual surrogate nodes, respectively. Consumers are denoted with \mathcal{D} , the set of time slots is represented with \mathcal{T} , while the set of stochastic scenarios is Φ . Each scenario $\phi \in \Phi$ has an associated realization probability, represented by p_ϕ . Consumers $d \in \mathcal{D}$ express a time-varying traffic demand for each scenario $\phi \in \Phi$, that we indicate with $r_d^{t,\phi}$. The CDN provider ensures that at least a fraction ϵ of the aggregate requests in every time slot is served by CDN nodes within a bounded delay, denoted as Δ . The topological information is encoded in our proposed optimization model using the $\delta_{s,d}$ input parameter, which represents the delay between client $d \in \mathcal{D}$ and CDN node $s \in \mathcal{S}$. \mathcal{K}_s^P and \mathcal{K}_s^V are the bandwidth capacities for physical and virtual CDNs, respectively. Physical nodes have an activation cost \mathcal{C}_s^P , while virtual CDN nodes have a traffic-proportional cost \mathcal{C}_s^V .

Our proposed optimization model chooses the optimal physical nodes placement and request routing. a_s is a binary decision variable that is set to 1 if and only if the physical candidate server $s \in \mathcal{S}^P$ is activated. Traffic requests for consumer d , in time slot t for scenario ϕ can be served by flows $y_{s,d}^{t,\phi}$ and $z_{s,d}^{t,\phi}$. In particular, $y_{s,d}^{t,\phi}$ is a flow originating from the physical node $s \in \mathcal{S}^P$, while $z_{s,d}^{t,\phi}$ is a flow provided by the virtual node $s \in \mathcal{S}^V$.

The deterministic equivalent program in the extensive form for the CDN planning problem of our infrastructure (EF-CDN) is formulated as follows:

$$\min \sum_{s \in \mathcal{S}^P} \left[\mathcal{C}_s^P a_s \right] + \mathbb{E}_\Phi \left[\sum_{s \in \mathcal{S}^V} \sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} (\mathcal{C}_s^V z_{s,d}^{t,\phi}) \right] \quad (1)$$

subject to:

$$\sum_{d \in \mathcal{D}} y_{s,d}^{t,\phi} \leq a_s \mathcal{K}_s^P \quad \forall s \in \mathcal{S}^P, t \in \mathcal{T}, \phi \in \Phi \quad (2)$$

$$\sum_{d \in \mathcal{D}} z_{s,d}^{t,\phi} \leq \mathcal{K}_s^V \quad \forall s \in \mathcal{S}^V, t \in \mathcal{T}, \phi \in \Phi \quad (3)$$

$$\sum_{s \in \mathcal{S}^P} y_{s,d}^{t,\phi} + \sum_{s \in \mathcal{S}^V} z_{s,d}^{t,\phi} = r_d^{t,\phi} \quad \forall d \in \mathcal{D}, t \in \mathcal{T}, \phi \in \Phi \quad (4)$$

$$\sum_{d \in \mathcal{D}} \left(\frac{\sum_{s \in \mathcal{S}^P | \delta_{s,d} \leq \Delta} y_{s,d}^{t,\phi} + \sum_{s \in \mathcal{S}^V | \delta_{s,d} \leq \Delta} z_{s,d}^{t,\phi}}{r_d^{t,\phi}} \right) \geq \epsilon \quad \forall t \in \mathcal{T}, \phi \in \Phi \quad (5)$$

$$a_s \in \{0, 1\} \quad \forall s \in \mathcal{S}^P \quad (6)$$

$$y_{s,d}^{t,\phi}, z_{s,d}^{t,\phi} \in \mathbb{R}^+ \quad \forall s \in \mathcal{S}, d \in \mathcal{D}, t \in \mathcal{T}, \phi \in \Phi. \quad (7)$$

The objective function (1) minimizes the overall costs given by the activation of physical CDN nodes as well as the usage of the virtual infrastructure. In particular, the virtual cost component is computed as the expected value for all the considered scenarios. Constraints (2) set a capacity bound on the overall demand served by physical CDN surrogates. If a physical surrogate is not activated, that is $a_s = 0$, it will not be capable to serve any request. Similarly, the virtual CDN nodes capacity is fixed in (3). In (4) we make sure that the overall clients' demands are served in any time slot and scenario, by virtual or physical surrogate servers. Flows can be split across multiple CDN servers. In (5) we control the overall service quality. We make sure that a fraction of at least ϵ requests in each time slot is served by CDN surrogates within a maximum delay of Δ . Finally, binary restrictions on the activation variables are set in (6), while non-negativity constraints on the continuous flow variables are enforced in (7). Rather than considering the worst or mean case, the stochastic formulation ensures that constraints hold in every scenario, while the objective function is optimized given the uncertainty on future traffic requests.

In order to solve the optimization problem (1)-(7) we employ different strategies:

- 1) A mixed integer linear programming solver (MILP);
- 2) The L-shaped algorithm (*single* and *multicut* versions);
- 3) A polynomial-time greedy heuristic.

In Sec. III-C we present our proposed heuristic, while we refer to Van Slyke and Wets [6] for an introduction to the L-shaped algorithm since, for the sake of brevity, in this paper we omit its implementation details.

C. Heuristic Solver

The L-shaped algorithm can be used to find an exact solution of the optimization problem we formulated in Sec. III-B. Theoretical results guarantee that the L-shaped algorithm converges to the optimal solution, but in some cases the speed of convergence might be too slow for the considered application. In this section we describe the polynomial-time

Algorithm 1: Heuristic solver

Input : $\langle \mathcal{D}, \mathcal{S}, \mathcal{T}, \Phi, r_d^{t,\phi}, \mathcal{K}_s^P, \mathcal{K}_s^V, \mathcal{C}_s^P, \mathcal{C}_s^V, \delta_{s,d}, \Delta, \epsilon, p_\phi \rangle$
Output: $\hat{a}, \hat{y}_{s,d}^{t,\phi}, \hat{z}_{s,d}^{t,\phi}, \min_cost$

- 1 $\hat{a} = [1; 1; \dots; 1]$; `phy_nodes` = `sort_phy_CDN_nodes()`;
- 2 **if** `¬is_feasible`(\hat{a}) **then**
| **return** `INFEASIBLE_ASSIGNMENT` ;
end
 $\langle \hat{y}_{s,d}^{t,\phi}, \hat{z}_{s,d}^{t,\phi}, \min_cost \rangle = \text{get_LP_solution}(\hat{a})$; `best_sol` = \hat{a} ;
- 3 **foreach** $s \in \text{phy_nodes}$ **do**
| $\hat{a}_s = 0$;
| $\langle \hat{y}_{s,d}^{t,\phi}, \hat{z}_{s,d}^{t,\phi}, \text{current_min_cost} \rangle = \text{get_LP_solution}(\hat{a})$;
| **if** `current_min_cost` $\geq \min_cost$ **then**
| | **break** ;
| **end**
| `min_cost` = `current_min_cost` ; `best_sol` = \hat{a} ;
end
- 4 $\hat{a} = \text{best_sol}$; $\langle \hat{y}_{s,d}^{t,\phi}, \hat{z}_{s,d}^{t,\phi}, \min_cost \rangle = \text{get_LP_solution}(\hat{a})$;

heuristic we designed to compute a close-to-optimal solution to the planning problem. Pseudo-code is provided in Algorithm 1 to describe the steps we use to achieve this purpose.

At the beginning of the algorithm, in Step 1, all the physical nodes are activated and sorted. In particular, we give higher priority to those nodes that can serve the largest amount of traffic demands within a delay of Δ . Given the structure of the problem, infeasible solutions are those that cannot be served even when all the physical CDN nodes have been activated. In Step 2 we check this condition and we eventually signal a potential infeasibility. The `get_LP_solution` function computes a solution for the continuous relaxation of the model (1)-(7), using a standard linear programming solver, and therefore has a polynomial-time complexity. As outputs, it returns the optimal flows for the physical and virtual CDN nodes as well as the overall cost. In case of infeasibility, the output value of `min_cost` is set to infinite. The loop in Step 3 deactivates at every iteration one new physical node, according to the previously generated ordering, and it completes when the objective function does not improve anymore. Lastly, in Step 4 we compute the optimal flows starting from the best physical nodes allocation choice.

IV. NUMERICAL RESULTS

In this section we present the numerical results we obtained performing a thorough analysis of our models and heuristics under realistic network conditions.

Unless stated otherwise, our network topology is created using the Barabási–Albert model and is composed of 50 consumers, 20 physical and 15 virtual CDN nodes. Traffic demands are generated using as a reference the Cisco VNI data for the 2014-2017 years: the planning horizon is of 3 years and we used 36 different time slots. Traffic uncertainty is taken into account by considering 10 different scenarios, with a variable overall demand uniformly distributed between 80% and 120% of the Cisco forecast. Slightly better results were observed considering the more favorable case of the binomial distribution with success probability $p = 0.5$. To control the overall demand, we limit to 20 Gbit/s the maximum traffic requests that a consumer can generate in a time slot of a scenario. Physical CDN nodes have a capacity of 12.5 Gbit/s, while virtual CDNs can serve up to 8 Gbit/s. Similar trends have been observed for other values of the capacity, omitted

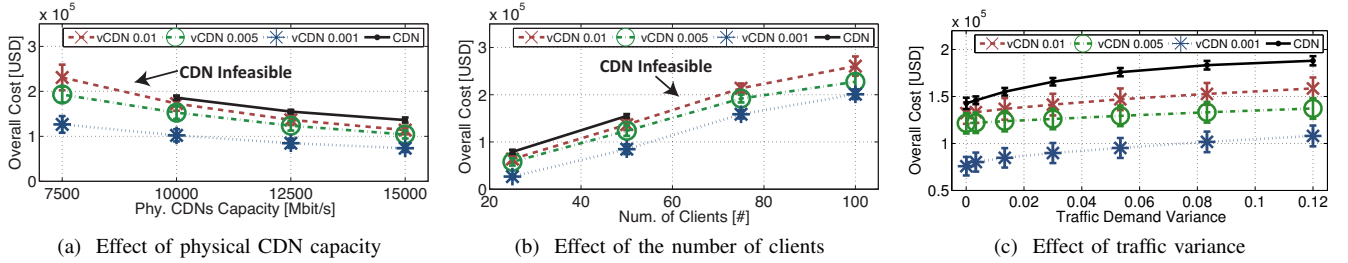


Figure 2. *NFV Benefits*. Plots 2a-2c show the cost benefits of an architecture composed of a mix of physical and virtual CDN nodes (denoted in the figures with vCDN), with respect to the scenario where only the physical CDN infrastructure is used (denoted in the figures with CDN). Different prices are considered for the vCDN case, as shown in the legend.

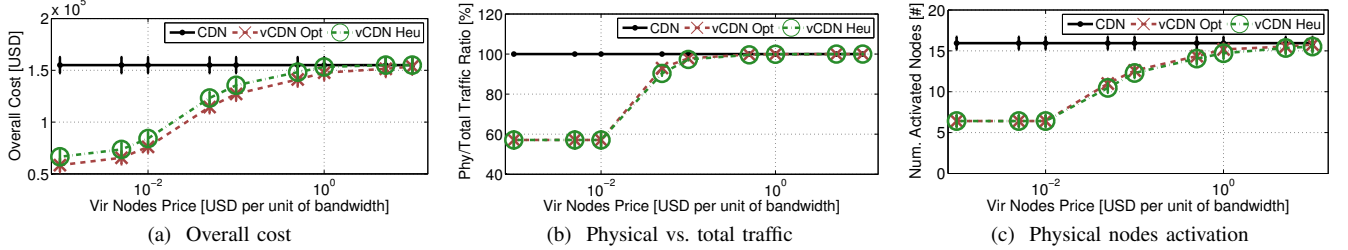


Figure 3. *Effect of the vCDN Price*. Plots 3a-3c show the effect of the virtual nodes pricing on the overall cost (Fig. 3a), the fraction of traffic served by the physical infrastructure (Fig. 3b) and the number of physical nodes activated (Fig. 3c), for different solution algorithms as well as a physical-only infrastructure (denoted in the figures with CDN).

here for the sake of brevity. Link delays are generated in the range of those available on Rocketfuel for the Sprintlink (US) topology, with an average delay of 3 ms. Moreover, we assume the CDN provider wants to guarantee that at least 95% of the requests are served by surrogates with a delay lower than 12 ms, that is, selected CDN nodes are, on average, 4 hops far from the consumer's location.

Lastly, prices are set as follows: we assume that the cost to install and operate one CDN node is set in the range [8; 12] kUSD, while different prices will be considered for virtual CDN nodes in the range [0.001; 10] USD per Mbit/s. For the same physical-virtual price ratios, even by considering different values for the physical CDN pricing, we observed similar trends as those discussed in this section. Hereafter, we report the result we obtained using CPLEX 12.5 as a MILP solver, bounding the maximum execution time of the algorithms to 1 hour (with a 5% MIP gap), and using a machine equipped with a quad-core Intel i7-3770 (3.40 GHz) CPU with 16 Gbyte of RAM. Lastly, for each of the results we performed 20 different runs and we report the narrow 95% confidence intervals.

NFV Benefits. Fig. 2 shows the cost benefits for the mixed physical-virtual CDN architecture, considering different virtual prices. As expected, in Fig. 2a-2c, lower CDN prices lead to lower overall costs. The effect of the physical CDN capacity is shown in Fig. 2a. If the capacity of the physical CDN appliances is lower than 10 Gbit/s, using virtual nodes becomes mandatory since otherwise an infeasibility is produced. When the physical CDN nodes capacity is set to 15 Gbit/s, cost savings up to 46% are experienced for cheap virtual CDN pricing (i.e., 0.001 USD per unit of bandwidth), whereas the saving is reduced to 16% if we set a vCDN price of one order of magnitude larger (i.e., 0.01 USD per unit of bandwidth). Fig. 2b shows the effect of the number of clients on the costs. The physical-only CDN infrastructure cannot handle

more than 50 clients, whereas up to 100 clients can be served if we also leverage the 15 virtual CDNs deployed. The effects of traffic uncertainties are quantified in Fig. 3c, where we show the overall cost as a function of the traffic demand variance. With the largest variance that we took into account, the vCDN infrastructure leads to cost savings in the range 16-43% according to the vCDN pricing.

Effect of the vCDN Price. Due to its remarkable effect, in Fig. 3 we show the impact of the virtual CDN pricing using the different solution algorithms. The overall cost is portrayed in Fig. 3a. Solutions obtained with exact solvers such as the deterministic equivalent program in the extensive form or the L-shaped algorithm lead to costs up to 11% (and on average 6%) lower than those reported with the heuristic. Fig. 3a clearly shows that the higher the prices, the lower the economic benefits of using a mixed physical-virtual CDN infrastructure. In particular, considering prices in the range [0.001; 0.5] USD per unit of bandwidth, the cost savings compared to the physical-only solution are in the range 5-64%. Fig. 3b shows the proportion of physical traffic with respect to the overall demand, as a function of the price of virtual nodes. Cheap prices make the virtual CDN capacity be fully saturated, and for this reason the left hand-side of Fig. 3b has an horizontal trend that accounts for 57% of the overall traffic. As a consequence, in Fig. 3c is reported that the number of physical CDN nodes does not increase for virtual prices lower than 0.01 USD. Comparing Fig. 3b and 3c, the slope of the curve is less steep in the second plot, since there are cases where it is convenient to strategically deploy a physical CDN server in a special position of the topology even though it is not fully used by the clients. It is interesting to note that for both Fig. 3b and 3c the heuristic algorithm leads to solutions that are practically overlapped to the optimal choice.

Computing time. To limit the effects of infeasibilities that negatively affects results on the computing time, hereafter we

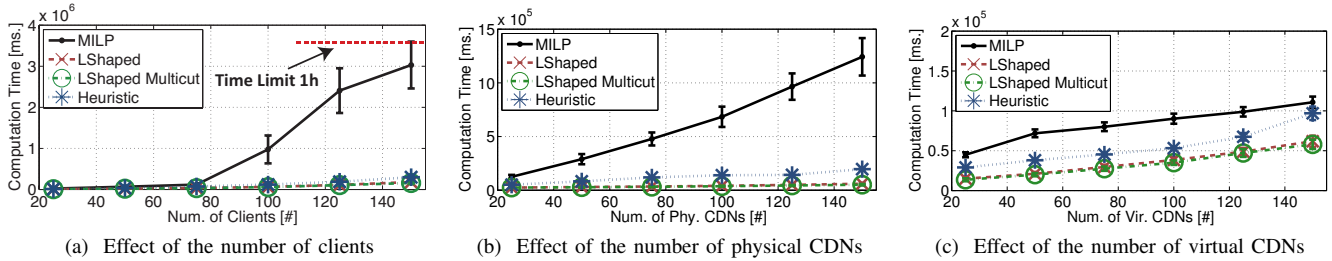


Figure 4. Execution time. Plots 4a–4c show the behavior of the different solution algorithms as a function of the number of clients as well as the number of surrogate nodes. We observe that the number of clients has the most remarkable effect on the execution time.

raise the number of virtual CDN nodes to 50, making the network be capable to serve up to 150 consumers. Fig. 4 shows the execution time of the different algorithms as a function of the number of clients (Fig. 4a), the number of physical (Fig. 4b), and virtual CDNs (Fig. 4c).

The number of clients is the parameter that mostly affects the execution time, as portrayed in Fig. 4a. In particular, the MILP solver for the deterministic equivalent program has a time trend that is exponential in the number of served clients (Fig. 4a), but linear in the number of physical (Fig. 4b) and virtual CDNs (Fig. 4c). While the MILP solver can hardly scale to topologies with a larger number of nodes, this possibility is instead offered by the L-shaped decomposition and our proposed heuristic. As a matter of fact, all these algorithms can solve the planning problem saving up to 94% of time compared to MILP, when considering 150 clients, as shown in Fig. 4a. Lastly, although there are cases where the heuristic algorithm is slightly slower than the L-shaped algorithm (as in Fig. 4c), we remark the fact that the heuristic has a worst-case polynomial time complexity, whereas a comparable theoretical result for the L-shaped algorithm does not hold.

V. CONCLUSION

In this paper we tackled the stochastic planning problem for content delivery to study potential benefits that Network Functions Virtualization can lead for content distribution purposes. We considered a mixed architecture where both physical as well as virtual CDN nodes can be used by a CDN owner to implement the content distribution service. The owner performs the planning choice for the physical CDN infrastructure on a long-term time schedule, possessing only a stochastic estimate of future traffic demands.

Our study shows that a mixed solution where both virtual and physical CDN nodes are used can dramatically reduce the overall costs sustained by the operator to purchase and operate the distribution infrastructure. In particular, we observed that gains can be up to 65% when considering the cheapest vCDN price. Our contribution is also to formulate efficient solution algorithms for the two-stage stochastic planning problem that can scale to realistic topology sizes. Rather than solving the deterministic equivalent problem in the extensive form, our proposed L-shaped algorithm and heuristic can efficiently find a solution, saving up to 95% of time compared to the MILP solver.

ACKNOWLEDGMENT

This work was partially supported by French ANR in the framework of the Green-Dyspan project.

REFERENCES

- [1] V. K. Adhikari, Y. Guo, F. Hao, M. Varvello, V. Hilt, M. Steiner, and Z.-L. Zhang, “Unreeling Netflix: Understanding and Improving Multi-CDN Movie Delivery,” in *Proc. of IEEE INFOCOM*. IEEE, Orlando, Florida, USA, Mar. 2012, pp. 1620–1628.
- [2] H. H. Liu, Y. Wang, Y. R. Yang, H. Wang, and C. Tian, “Optimizing Cost and Performance for Content Multihoming,” in *Proc. of ACM SIGCOMM*. ACM, Helsinki, Finland, Aug. 2012, pp. 371–382.
- [3] E. Nygren, R. K. Sitaraman, and J. Sun, “The Akamai Network: A Platform for High-Performance Internet Applications,” *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 2–19, July 2010.
- [4] B. Niven-Jenkins, F. L. Faucheur, and N. Bitar, “Content Distribution Network Interconnection (CDNI) Problem Statement,” RFC 6707 (Informational), Internet Engineering Task Force, Sep. 2012.
- [5] ETSI, “Network Functions Virtualisation - White Paper,” http://portal.etsi.org/NFV/NFV_White_Paper.pdf, Last accessed: May 2014.
- [6] R. M. Van Slyke and R. Wets, “L-shaped linear programs with applications to optimal control and stochastic programming,” *SIAM Journal on Applied Mathematics*, vol. 17, no. 4, pp. 638–663, 1969.
- [7] N. M. K. Chowdhury and R. Boutaba, “Network Virtualization: State of the Art and Research Challenges,” *IEEE Communications Magazine*, vol. 47, no. 7, pp. 20–26, 2009.
- [8] A. Fischer, J. Botero, M. Till Beck, H. de Meer, and X. Hesselbach, “Virtual Network Embedding: A Survey,” *IEEE Communications Surveys Tutorials*, vol. 15, no. 4, pp. 1888–1906, 2013.
- [9] A. Jarray and A. Karmouch, “Decomposition Approaches for Virtual Network Embedding With One-Shot Node and Link Mapping,” *IEEE/ACM Transactions on Networking (TON)*, accepted for publication, March 2014.
- [10] N. M. K. Chowdhury, M. R. Rahman, and R. Boutaba, “Virtual network embedding with coordinated node and link mapping,” in *Proc. of IEEE INFOCOM*. IEEE, Rio de Janeiro, Brazil, Apr. 2009, pp. 783–791.
- [11] X. Cheng, S. Su, Z. Zhang, H. Wang, F. Yang, Y. Luo, and J. Wang, “Virtual Network Embedding Through Topology-Aware Node Ranking,” *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 2, pp. 38–47, 2011.
- [12] J. M. Wang, J. Zhang, and B. Brahim, “Content Multi-homing: an Alternative Approach,” in *Proc. of the IEEE Int. Conf. on Communications (ICC)*. IEEE, Sydney, Australia, Jun. 2014, pp. 371–382.
- [13] A. Atamtürk and M. Zhang, “Two-Stage Robust Network Flow and Design Under Demand Uncertainty,” *Operations Research*, vol. 55, no. 4, pp. 662–673, 2007.
- [14] X. Liu, “The Role of Stochastic Programming in Communication Network Design,” *Computers & Operations Research*, vol. 32, no. 9, pp. 2329–2349, 2005.
- [15] A. Eisenblätter and J. Schweiger, “Multistage Stochastic Programming in Strategic Telecommunication Network Planning,” *Computational Management Science*, vol. 9, no. 3, pp. 303–321, 2012.
- [16] F. Figueiredo, F. Benevenuto, and J. M. Almeida, “The Tube over Time: Characterizing Popularity Growth of YouTube Videos,” in *Proc. of the 4th Int. Conf. on Web Search and Data Mining (WSDM)*. ACM, Hong Kong, Feb. 2011, pp. 745–754.