

Content-Aware Planning Models for Information-Centric Networking

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

^{*} LRI, Université Paris-Sud

[†] DEIB, Politecnico di Milano

{michele.mangili, fabio.martignon}@lri.fr {capone, malucelli}@elet.polimi.it

[‡] IUF, Institut Universitaire de France

Abstract—Information-Centric Networking (ICN) has recently gained momentum as a promising paradigm for the next-generation Internet architecture. The first prototypes for ICN-capable routers have already been developed, and network operators will soon have the opportunity to experience the advantages introduced by this technology. However, to migrate the devices to this novel architecture, non-negligible investments should be made. Therefore, it is of utter importance to provide clear quantitative insights of the expected economic benefits that operators will experience by switching to the ICN paradigm. For these reasons, in this paper we tackle the *content-aware network-planning* problem, and we formulate a novel optimization model to study the migration to an ICN, in a budget-constrained scenario. Our formulation takes into accurate account 1) traffic routing and 2) content caching. We further complement our contribution by designing a Randomized Rounding heuristic that scales up to realistic topologies composed of hundreds of nodes.

Keywords—Information-Centric Networking, Planning, Optimization, Content-Distribution.

I. INTRODUCTION

Recent traffic measurements have clearly shown that more than 50% of the overall Internet traffic is generated to retrieve contents, as illustrated in the Sandvine Report and Cisco VNI forecasts [1], [2]. However, being able to accommodate the content distribution needs of the users is still in today’s Internet a challenging task, and adequate technical solutions such as Content Delivery Networks (CDNs) have specifically been designed to achieve this objective [3].

Meanwhile, innovative paradigms known under the name of Information-Centric Networking (ICN) have recently gained momentum in the research community. Despite the fact that there are many different designs that belong to this category, all of them are based on the idea that by directly intervening on the protocol stack, the content distribution capabilities of the network can be boosted [4]. Among all the advantages that can be experienced by switching to this novel infrastructure, *traffic offloading* stems out as being the most relevant achievement [5]. Despite that, other advantages can also be gained: lower delays, better security and multipath routing all integrate as positive features of these networks [6].

Rather than being in their preliminary steps, these research efforts have already reached the point where the first working prototypes for ICN-enabled routers have been realized by Alcatel [7], Cisco [8] and Parc [9]. Specific hardware and software components are required in order to support ICN

packet forwarding at wire-speed, and thus operators will certainly have to make non-negligible investments to purchase the new ICN devices. As a result, they will be willing to transition their infrastructures to ICN only if clear economic benefits are envisioned: by upgrading a router to ICN and by installing a given amount of storage to memorize frequently requested contents, the router will behave as a distributed cache. In this way, the operator can experience significant economic savings accountable to traffic offloading.

To pave the way for a potential paradigm shift from a TCP/IP network to ICN, we specifically consider the *migration* step to the ICN architecture and we formulate a novel *content-aware network planning model* that we use to compute the optimal *migration strategy* for the operator. On top of that, by considering relevant economic parameters, our model can also be used to understand which economic benefits are expected as a result of the transition to ICN. To achieve all these objectives, we take into account three economic parameters: 1) a traffic-proportional link cost, 2) the router migration cost and 3) the storage cost, proportional to the amount of memory installed at a given ICN-migrated node.

To summarize, in this paper we provide the following contributions:

- 1) We formulate a model to evaluate the optimal content-distribution performance of an IP network under unsplitable routing conditions.
- 2) We propose a novel *content-aware network-planning* Mixed Integer Linear Programming (MILP) model for the *migration* to an ICN. Our formulation determines the optimal node migration and cache allocation in a budget-constrained scenario. Unsplittable routing conditions are still enforced by non-migrated routers.
- 3) We design a near-optimal Randomized Rounding heuristic that is capable to scale up to realistic network sizes composed of hundreds of nodes.
- 4) We quantitatively evaluate the benefits of migrating to an ICN, with different pricing configurations.

Our key findings suggest that 1) by migrating only few nodes to ICN, remarkable traffic reductions will be experienced by the operator; 2) ICN benefits also content providers since it significantly offloads their distribution infrastructures, and 3) when the content popularity distribution is very skewed, the storage space installed at the migrated nodes is an order of magnitude smaller than for less skewed distributions. To

the best of our knowledge this is the first paper that tackles the content-aware network-planning problem for ICN, by explicitly taking into account the migration, storage and traffic costs.

This paper is structured as follows: related works are discussed in Sec. II, while in Sec. III we introduce the system model. In Sec. IV we extensively describe the optimization models we use to compute the overall content delivery cost of an IP network and the content-aware planning model for migration to an ICN. In Sec. V we illustrate the randomized rounding heuristic for ICN, while numerical results are discussed in Sec. VI. Finally, concluding remarks are presented in Sec. VII.

II. RELATED WORK

ICN-capable routers are beginning to appear, and some prototypes, peaking the remarkable throughput of 12 Tbps [9], have already been presented by Alcatel [7], Cisco [8] and Parc [9]. However, the design of these devices demands for specific hardware and software solutions to make them operate at wire speed, and these strict requirements will likely have remarkable effects on the pricing of the equipment.

A first investigation on the possible architecture of an ICN router, with special attention towards computational issues related to the content store, has been originally provided by Arianfar et al. in [10]. Perino et al. have further complemented such analysis by presenting in [11] clear quantitative insights on the memory technologies that can be used to make wire-speed processing of ICN packets a reality. In both these works, preliminary economic data especially related to the prices of memories have been provided.

Another branch of research is devoted to optimal cache placement and request routing for content dissemination in the network. To the best of our knowledge, we are the first to tackle the planning problem to specifically study the migration to an ICN, while previous works have mostly addressed performance optimization for Internet content distribution.

A pioneering work by Krishnan et al. is presented in [12], and deals with cache placement in a TCP/IP network to minimize the overall network flow. Among the key-features of their formulation, we mention that they bound the number of caches that can be installed, moreover they assume the average *flow hit-rate* is given as an input parameter. Wang et al. formulate in [13] a model to solve a storage constrained cache allocation problem with optimal object placement in ICN. They focus the analysis on discovering which parameters mostly affect the location of caches in the topology. In [14], Hasan et al. tackle the problem of minimizing the overall cost for inter-Autonomous System cache deployments in transit ISP networks, considering the server, energy and bandwidth prices. Finally, optimal content-oriented request routing is investigated by Mihara et al. in [15]. They minimize the overall traffic on the most congested link, however caching is not considered in their analytical framework.

Our MILP formulation differs from previous works for the following reasons: 1) we accurately model link capacities and traffic flows, 2) we explicitly take into account the contents (i.e., the objects), 3) we adopt an economic perspective on

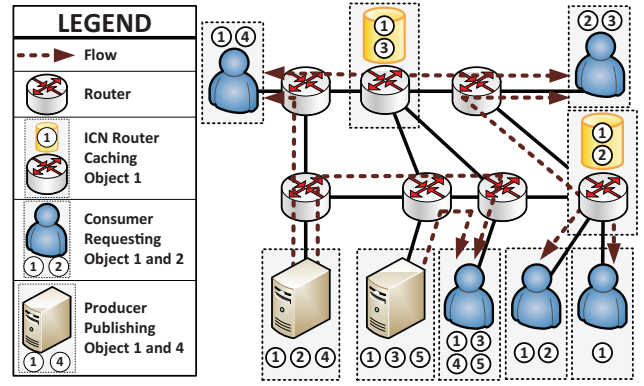


Figure 1. *System model.* Given the network topology, consumers’ requests, and objects served by content providers, our optimization model chooses which routers should be migrated to ICN, and which objects should they cache. The ICN is an overlay built on top of the underlying network connectivity.

the subject, solving the network planning problem for the migration to an ICN and 4) we jointly solve the optimal *request routing*, *cache provisioning* and *object placement* problems in a budget-constrained scenario.

III. SYSTEM MODEL

In this section we describe the system model and we discuss the rationale of our approach. A comprehensive introduction to some of the most notable ICN proposals can be found in [5].

Fig. 1 represents an example to describe relevant features of our proposed system model. Three types of nodes are available in the topology: consumers, producers and routers. All the nodes operate on a finite set of contents, called the “*catalog*”. For the sake of simplicity, as depicted in Fig. 1, in this example we assume that the catalog is composed of 5 objects. Producers publish objects in the network, whereas consumers generate demands for them. It is possible that the same object is provided by different producers, as represented in the figure.

Each link in the network is characterized by having a traffic-proportional price and a bounded capacity. ICN functionalities are built as an overlay on top of a TCP/IP network. Since ICN routers can provide in-network caching functionalities, the operator can significantly reduce its traffic costs, by migrating routers to this paradigm and by moving content replicas closer to the location where most of the users are requesting them. However, to perform the migration, the operator must pay the corresponding costs which are given by 1) the price to migrate a router to ICN, C^M and 2) the storage price to memorize one object at a migrated router, C^S . We bound the migration costs (i.e., those due to node migration and caching storage) to the value B , that is the total migration budget the operator is willing to spend. Being implemented as an overlay, ICN routers issue upstream traffic requests as if they were the request origins, finally, on top of offering caching functionalities, they also support splittable request routing.

In the rest of the paper we describe our proposed optimization models and heuristic to help the operator determine the optimal ICN node migration, object placement and request routing.

Table I. SUMMARY OF THE NOTATION USED IN THIS PAPER.

Parameters of the Models	
A	Set of arcs
N, C, P, R	N Set of nodes, $C \subset N$ Set of consumers, $P \subset N$ Set of producers, $R \subset N$ Set of routers
Q	Set of requesters. In the IP network model $Q = C$. In the ICN model $Q = C \cup R$.
O	Set of objects
$FS(i)$	Set of forward arcs $(i, j) \in A$ for node $i \in N$
$BS(i)$	Set of backward arcs $(j, i) \in A$ for node $i \in N$
$b_{i,j}$	Capacity of arc $(i, j) \in A$
$p_{i,j}$	Price per unit of traffic on arc $(i, j) \in A$
d_c^o	Demand of consumer $c \in C$ for object $o \in O$
r_p^o	0-1 Object reachability matrix $r_p^o = 1$ if producer $p \in P$ can serve object $o \in O$
C^M, C^S	Price to migrate one router to ICN, C^M Price to install one storage unit, C^S
B	Total migration budget

Decision Variables of the Models	
$y_{i,j}^{o,q}$	Flow on arc $(i, j) \in A$ for object $o \in O$, requested by requester $q \in Q$
$z_{i,j}^q$	0-1 Routing variable: $z_{i,j}^q = 1$ if arc $(i, j) \in A$ can be used to route requests for requester $q \in Q$
m_r	0-1 Router migration variable $m_r = 1$ if router $r \in R$ migrates to ICN.
k_r^o	0-1 Cache storage variable $k_r^o = 1$ if ICN router $r \in R$ caches object $o \in O$
w_l^o	Flow served by producer or router $l \in (P \cup R)$ for object $o \in O$, when l stores a replica of o
$F_r^{o,q}$	Flow balance at router $r \in R$, for object $o \in O$, requested by $q \in Q$

IV. DESIGN MODELS

In this section we describe the optimization models we use to evaluate the migration to an ICN. Sec. IV-A presents the IP network model, while Sec. IV-B is devoted to the ICN network planning formulation.

Let us introduce the notation used in describing the planning problems and in the optimization models. We represent the network as a directed graph $G = (N, A)$, where the set of nodes N is partitioned into consumers C , producers P , and routers R (i.e.: $N = C \cup P \cup R$). The set of forward and backward arcs of node $i \in N$ are denoted with $BS(i)$ and $FS(i)$, respectively. Since in the planning we consider only the downstream, producers have no incoming arcs and consumers have no outgoing arcs. Network arcs $(i, j) \in A$ are characterized by a capacity, denoted with $b_{i,j}$, and a price per unit of traffic, $p_{i,j}$. We denote with O the set of objects, and we assume that all of them have the same size, as frequently done in the literature (e.g.: [13], [16]). Let Q be the set of requesters; for both the IP and ICN network models, requesters are nodes from which traffic requests originate: in the IP network, only the consumers can behave as such, and thus $Q \equiv C$. Each consumer $c \in C$ expresses a traffic demand d_c^o for object $o \in O$. Producers can serve a subset of the entire object catalog, in particular we represent with the binary parameter r_p^o the object reachability matrix ($r_p^o = 1$ if producer $p \in P$ publishes object $o \in O$, otherwise $r_p^o = 0$). For the sake of clarity, in Table I, we summarize the notation used throughout the paper.

A. IP Network model

We start describing the IP network model we use as a benchmark with respect to the solution we get when studying the planning of an ICN. In the IP routing problem, objects must be routed from producers where they are available to consumers, possibly passing through routers, at the minimum overall cost. We assume that flows are unsplittable.

The problem can be naturally described as a multicommodity flow model, where a commodity is associated with every pair {object, requester}. Let variables $y_{i,j}^{o,q}$ denote the flow of object $o \in O$ on arc $(i, j) \in A$ for requester $q \in Q$. In addition, in order to account for the unsplittable flow requirement we introduce binary variables $z_{i,j}^q$ whose value is 1 if arc $(i, j) \in A$ is used to route traffic for requester $q \in Q$, and 0, otherwise. Note that another aspect of the problem involves the selection of the producer to serve each request, in the presence of multiple copies of some objects. To account for that, variables w_p^o denoting the actual quantity of flow of object o referring to producer p must be introduced.

The minimum cost request routing problem under unsplittable flow conditions for an IP network can therefore be formulated as follows:

$$\min \sum_{(i,j) \in A} p_{i,j} \sum_{o \in O} \sum_{q \in Q} y_{i,j}^{o,q} \quad (1)$$

subject to:

$$\sum_{(j,r) \in BS(r)} y_{j,r}^{o,q} - \sum_{(r,j) \in FS(r)} y_{r,j}^{o,q} = 0 \quad \forall o \in O, \forall q \in Q, \forall r \in R \quad (2)$$

$$\sum_{(j,i) \in BS(i)} y_{j,i}^{o,i} = d_i^o \quad \forall o \in O, \forall i \in C \quad (3)$$

$$\sum_{q \in Q} \sum_{(p,j) \in FS(p)} y_{p,j}^{o,q} = w_p^o \quad \forall o \in O, \forall p \in P \quad (4)$$

$$w_p^o \leq r_p^o \sum_{c \in C} d_c^o \quad \forall p \in P, \forall o \in O \quad (5)$$

$$\sum_{c \in C} d_c^o = \sum_{p \in P} w_p^o \quad \forall o \in O \quad (6)$$

$$\sum_{o \in O} \sum_{q \in Q} y_{i,j}^{o,q} \leq b_{i,j} \quad \forall (i, j) \in A \quad (7)$$

$$\sum_{o \in O} y_{i,j}^{o,q} \leq b_{i,j} z_{i,j}^q \quad \forall i \in N \setminus C, \forall (i, j) \in FS(i), \forall q \in Q \quad (8)$$

$$\sum_{(i,n) \in FS(i)} z_{i,n}^q \leq 1 \quad \forall i \in N \setminus C, \forall q \in Q \quad (9)$$

$$z_{i,j}^q \in \{0, 1\} \quad \forall q \in Q, \forall (i, j) \in A \quad (10)$$

$$w_p^o \geq 0 \quad \forall p \in P, \forall o \in O \quad (11)$$

$$y_{i,j}^{o,q} \geq 0 \quad \forall o \in O, \forall q \in Q, \forall (i, j) \in A. \quad (12)$$

The objective function (1) minimizes the overall traffic costs incurred by the provider on all network arcs.

The flow balance at every router and consumer node is imposed by (2) and (3), respectively. The balance at producer nodes depends on the entering flow of each object (4) which is regulated by (5) and (6). These constraints consider the fact that requests can be served only by those producers that are actually publishing a copy of the given object in the network, and that the overall traffic served by the producers equals the overall demand expressed by the consumers.

Capacity constraints are enforced in (7). Unsplittable routing conditions are imposed in (8) and (9). In particular, the set of constraints (8) makes sure that flows for requester $q \in Q$ are forwarded only on the arcs $(i, j) \in A$ where $z_{i,j}^q = 1$, whereas in (9) we make sure that routers and producers have at most only one egress arc for requester $q \in Q$.

Finally, non negativity on flow variables and binary condition on $z_{i,j}^q$ are imposed in (10)-(12). Notice that if 0-1 variables $z_{i,j}^q$ are fixed, the problem amounts to a continuous multicommodity flow that can be solved by standard linear programming solvers.

B. ICN Planning

We now extend the model presented in Sec. IV-A to solve the *content-aware network planning* problem in ICN.

Let C^M denote the additional cost to migrate one IP router to ICN. Once that a router has been migrated to this paradigm, caching storage can be installed on it. C^S denotes the storage cost to add the caching space sufficient to memorize one object. The overall *migration cost* should not exceed the total available budget, which is denoted with B . Two sets of binary variables are used in the ICN planning model: m_r and k_r^o . They are such that $m_r = 1$ if router $r \in R$ migrates to ICN, otherwise $m_r = 0$; similarly $k_r^o = 1$ if router $r \in R$ caches object $o \in O$, while $k_r^o = 0$ if the object is not cached.

Given the above definitions we formulate the budget-constrained ICN planning problem as follows:

$$\min \sum_{(i,j) \in A} p_{i,j} \sum_{\substack{o \in O \\ q \in Q}} y_{i,j}^{o,q} + \left[C^M \sum_{r \in R} m_r + C^S \sum_{r \in R} \sum_{o \in O} k_r^o \right] \quad (13)$$

subject to (3)-(5), (7)-(8), (10), and (12)

$$\sum_{(j,r) \in BS(r)} y_{j,r}^{o,q} - \sum_{(r,j) \in FS(r)} y_{r,j}^{o,q} = F_r^{o,q} \quad \forall o \in O, \forall q \in Q, \forall r \in R \quad (14)$$

$$- \sum_{q \in Q} F_r^{o,q} = w_r^o + \sum_{(i,r) \in BS(r)} y_{i,r}^{o,r} \quad \forall o \in O, \forall r \in R \quad (15)$$

$$w_r^o \leq k_r^o \cdot \sum_{c \in C} d_c^o \quad \forall o \in O, \forall r \in R \quad (16)$$

$$k_r^o \leq m_r \quad \forall o \in O, \forall r \in R \quad (17)$$

$$y_{i,j}^{o,r} \leq m_r b_{i,j} \quad \forall (i,j) \in A, \forall o \in O, \forall r \in R \quad (18)$$

$$\sum_{(i,j) \in FS(i)} z_{i,j}^q \leq 1 \quad \forall i \in P, \forall q \in Q \quad (19)$$

$$\sum_{(r,j) \in FS(r)} z_{r,j}^q \leq 1 + m_r \cdot (|FS(r)| - 1) \quad \forall r \in R, \forall q \in Q \quad (20)$$

$$\sum_{c \in C} d_c^o = \sum_{l \in P \cup R} w_l^o \quad \forall o \in O \quad (21)$$

$$C^M \sum_{r \in R} m_r + C^S \sum_{r \in R} \sum_{o \in O} k_r^o \leq B \quad (22)$$

$$y_{r,i}^{o,r} = 0 \quad \forall r \in R, \forall (r,i) \in FS(r), \forall o \in O \quad (23)$$

$$w_l^o \geq 0 \quad \forall l \in P \cup R, \forall o \in O. \quad (24)$$

The objective function (13) takes into account traffic and migration cost components. The former is given by $\sum_{(i,j) \in A} p_{i,j} \sum_{\substack{o \in O \\ q \in Q}} y_{i,j}^{o,q}$, the latter is instead the sum of node migration costs $C^M \sum_{r \in R} m_r$, and storage costs $C^S \sum_{r \in R} \sum_{o \in O} k_r^o$.

Flow balance constraints for routers are expressed in (14). In particular, if a router $r \in R$ migrates to ICN (i.e. $m_r = 1$), we let the flow balance be $F_r^{o,q} \leq 0$, meaning that r can directly serve incoming requests; otherwise, if $m_r = 0$ we set $F_r^{o,q} = 0$. The set of constraints (15) permits a router $r \in R$ to have caching functionalities (i.e. $w_r^o \geq 0$); furthermore it lets r behave as a requester (i.e. $y_{i,r}^{o,r} \geq 0$), a feature that facilitates traffic splitting in the network. The joint presence of constraints (16)-(18) makes sure that only ICN-migrated routers can provide caching functionalities and behave as requesters. In-network caching features are modeled in (16) and (17). In particular, if a router r migrates to ICN and stores in its local cache a copy of object o , it is then capable of directly serving upstream requests for that particular object. Instead, in (18) we prevent non-migrated routers behave as requesters.

Unsplittable request routing is enforced in the set of constraints (19) (*for producers only*) and (20) (*for routers only*). In particular, this latter set of constraints lets a migrated ICN router $r \in R$ to support splittable routing: if the router does not migrate to ICN (i.e. $m_r = 0$), then at most one egress link is used to route requests for $q \in Q$, otherwise if $m_r = 1$, then all the egress links can be used making the ICN-migrated router capable to perform multipath routing. In (21), we impose the condition that the overall demand generated for object $o \in O$ is satisfied by producers and caching routers. The budget allocated for the migration is limited by (22). The set of constraints (23) avoids loops, preventing requests expressed by a router to be fulfilled by the router itself, while in (24) non-negativity on flow variables is enforced.

V. RANDOMIZED ROUNDING HEURISTIC

As we will discuss in Sec. VI, even by using best of breed ILP solvers available today, the optimal solution of the ICN model (formulated in Sec. IV-B) can hardly be computed for very large topologies. For this reason, in this section we propose a Randomized Rounding (RR) heuristic for the ICN planning problem to efficiently compute the node migration and object allocation.

Algorithm 1: Randomized Rounding

Input : $model \leftarrow \langle b_{i,j}, p_{i,j}, d_c^o, r_p^o, C^M, C^S, B \rangle$
Output: obj_fun

- 1 $\hat{k}_r^o \leftarrow \text{SolveRelaxedICNModel}(model)$;
- 2 $\hat{k}_r^{max} \leftarrow \max_{o \in O} \hat{k}_r^o$;
- 3 $RL \leftarrow \text{SortRoutersByCumulativeProbabilityPerObject}(\hat{k}_r^o)$;
 $C \leftarrow 0$;
- foreach** $r \in RL$ **do**
 - $\bar{m}_r \leftarrow false$;
 - foreach** $o \in O$ **do**
 - 4 $w \leftarrow \{ \text{UniformRndValue}(0, 1) \leq (\hat{k}_r^o / \hat{k}_r^{max}) \}$;
 - 5 **if** $w \wedge (C < B)$ **then**
 - 6 **if** $\neg \bar{m}_r$ **then** $C \leftarrow C + C^M$
 - 7 $C \leftarrow C + C^S$; $\bar{k}_r^o \leftarrow true$; $\bar{m}_r \leftarrow true$;
 - end**
 - end**
- 8 $obj_fun \leftarrow \text{SolveICNModel}(mdl, \bar{k}_r^o, \bar{m}_r)$;

Algorithm 1 illustrates the heuristic in pseudo-code. The rationale behind it is to solve the continuous relaxation of the ICN model described in Sec. IV-B, computing the optimum fractional values of \hat{k}_r^o . We then interpret \hat{k}_r^o as the probability that object $o \in O$ is placed in the cache of the migrated router $r \in R$. As frequently done in the randomized rounding literature [17], we scale the relaxed variables for object caching \hat{k}_r^o dividing them by \hat{k}_r^{max} . The scaling is done to increase the object caching probability. Then, we assign a value to the suboptimal binary variables \bar{k}_r^o , setting them to one with a probability equal to \hat{k}_r^o . As a result, the algorithm chooses the node migration \bar{m}_r and object caching variables \bar{k}_r^o .

The solution of the continuous relaxation of the ICN model is computed in Step 1 of Alg. 1. In Step 2, the algorithm extracts \hat{k}_r^{max} , that is the largest \hat{k}_r^o value for each router. Instead, the cumulative caching probability for all the objects (i.e, the value $\sum_{o \in O} \hat{k}_r^o$) is computed in Step 3, where we also sort the routers in non-increasing order according to such metric. The overall migration costs are denoted by C . In Steps 4 and 5, the randomized caching choice is performed: the algorithm caches $o \in O$ at router $r \in R$ with probability $\hat{k}_r^o / \hat{k}_r^{max}$ if and only if there exists sufficient spare budget. In Step 6 the node migration costs are added to the value of C , while in Step 7, the storage costs are included and the caching variables are set. Finally, in Step 8, we solve the ICN model by fixing the caching and migration variables.

VI. NUMERICAL RESULTS

In this section we present the numerical results obtained performing extensive analysis using our *content-aware network planning models* and the corresponding *RR heuristic*.

Five topologies have been considered: Netrail (7 nodes), Abilene (11 nodes), Claranet (15 nodes), Airtel (16 nodes) and Géant (27 nodes) [18]. We uniformly distribute 5 producers and 10 consumers in the network, connecting them to at most one router. All network links have a capacity of 10 Gbit/s, and each consumer generates an aggregate demand of 1 Gbit/s distributed on the object catalog according to the Zipf popularity distribution. Two Zipf alpha exponents have

been considered: $\alpha = 1.2$ is used to model a very skewed popularity distribution where few objects are frequently requested, whereas $\alpha = 0.8$ better represents less skewed demands. The object catalog is composed of 10^8 different packet *chunks* of 4kB each, as in [19], [20]. For scalability reasons, and as frequently done in the ICN literature [19], we aggregate the traffic demands on 100 popularity classes; in other words, we solve the planning problem setting $|O| = 100$. We further assume that traffic demands are expressed by the users for a mid-term timespan of one year, and thus 37 Pbytes will be transferred by the network to the consumers.

To transfer 1 Gbyte of data, Amazon nowadays charges a variable price in the range [0.05; 0.12] USD. Given such pricing, if 1 Gbit/s is constantly transferred on a link, its yearly cost will be in the range [197k; 473k] USD; therefore, we uniformly generate the link price values ($p_{i,j}$) accordingly. Let $\max_p = 473k$ USD be the maximum yearly cost that the operator has to pay to satisfy the consumer's demand. We assume the cost to install one unit of storage is equivalent to 1/100 of the yearly traffic cost, (i.e, $C^S = 0.01 \max_p$), and similarly we set $C^M = \max_p$ for the router migration. Finally, we assume that the total migration budget B is in the range $B \in [1; 7] \max_p$, and therefore we let at most 7 nodes migrate. For each analysis, we performed 20 different runs and we computed the 95% confidence intervals depicted in the figures. For the sake of brevity, in this paper we present the most remarkable results, while the full set of plots is available online [18].

Computing time. Table II reports the average computing time necessary to solve different network instance of our planning problems using the CPLEX 12.5 solver on a Dual Intel Xeon E5-2630 v2 @ 2.60GHz machine with 64 GByte of RAM. The table refers to the scenario with budget $B = 7 \max_p$. For the optimal ICN model (OPT), we set the relative MIP gap tolerance between the best integer objective and the objective of the best LP relaxation to 1%. As shown in Table II, the solution of OPT is strongly dependent on the value of α and $|N|$: when $\alpha = 0.8$, CPLEX does not reach in 1 hour the MIP gap tolerance of 1% even in the Netrail topology, while 8.6 seconds are sufficient for $\alpha = 1.2$. On the other hand, the completion time of the RR heuristic is independent of α , and increases much less than OPT as $|N|$ increases, making it possible to use the RR algorithm even with realistic network topologies composed of hundreds of nodes.

Example Scenario. Fig. 2b represents the solution we observed while considering the Abilene network topology for the IP network model. Despite the fact that this result refers to a Zipf $\alpha = 0.8$, producers have a remarkably different load; in particular the first and the second most popular objects are

Table II. AVERAGE EXECUTION TIME: OPTIMIZATION MODEL (OPT) AND RANDOMIZED ROUNDING (RR)

Topology ($ N $)	OPT	OPT	RR	RR
	$\alpha = 0.8$	$\alpha = 1.2$	$\alpha = 0.8$	$\alpha = 1.2$
Netrail (22)	> 60min	8.6 s.	27.4 s.	28.8 s.
Abilene (26)	> 60min	41.5 s.	46.4 s.	46.0s.
Claranet (30)	> 60min	42.5 s.	60.6 s.	57.5 s.
Airtel (31)	> 60min	100.1 s.	77.1 s.	75.8 s.
Géant (42)	> 60min	675.7 s.	155.5 s.	153.4 s.

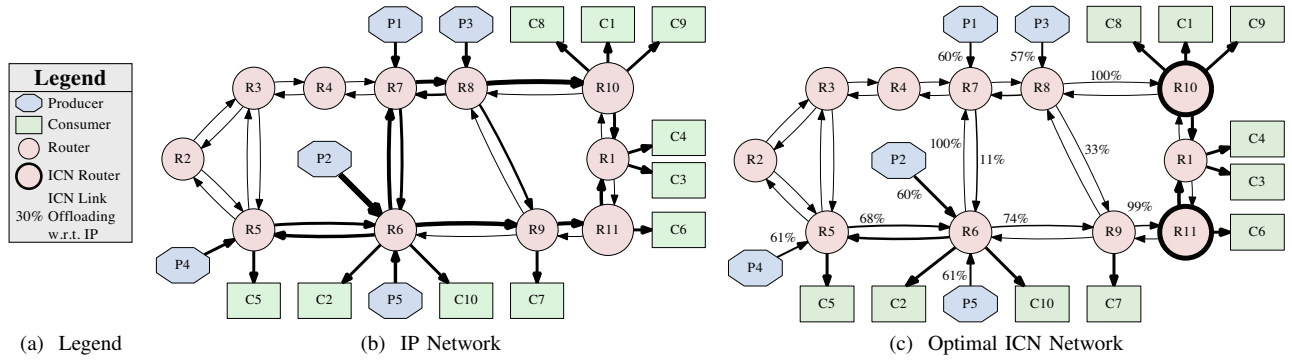


Figure 2. *Example Scenario.* We consider the Abilene topology with $\alpha = 0.8$ and budget of $B = 3.5 \max_p$. Fig. 2b represents the IP network, while Fig. 2c is the solution of the optimal ICN planning. The overall cost for the IP network is $14.3 \cdot 10^6$ USD, while in the optimal ICN network it decreases to $7.6 \cdot 10^6$ USD. Line size is proportional to the aggregate traffic the link is transferring, while the link offloading percentage reported in Fig. 2c compares the traffic in ICN w.r.t. the one of the IP network.

published by producer P2 and P5, respectively, thus their links are the most congested.

By comparing the solution depicted in Fig. 2b and Fig. 2c, we observe that migrating to ICN leads to an overall cost of $7.6 \cdot 10^6$ USD, compared to $14.3 \cdot 10^6$ USD for the IP network. Furthermore, such migration also reduces the link and producer congestion: in fact, on average, producers in ICN are providing 40% less traffic than in the IP solution. In general, network links are much less congested thanks to the presence of the two caches installed at router R10 and R11. In addition, there exist some arcs, such as the one between R6 and R7, that are not carrying traffic anymore. Another interesting observation is that while router R10 is the one that is serving the highest number of consumers, router R11 is preferred by the model over R1. Therefore, the best network planning strategy cannot take into account only the number of consumers a router is serving, but it requires an adequate planning model such as the one proposed in this paper.

Effect of the Budget. Fig. 3b-3i show the effect of the budget for both the Abilene and the Géant topology. The horizontal line represents the reference value of the IP model, where no router can migrate to ICN. The number of migrated nodes in the Abilene topology with $\alpha = 0.8$ or 1.2 is shown in Fig. 3b and 3c. For both scenarios, the RR heuristic deploys more ICN routers, especially when the available budget is large; in particular, on average, it migrates 11% more routers than the optimal solution, for $\alpha = 1.2$. In the Abilene network, at most 4 nodes are migrated to ICN, and the larger the α , the higher the number of migrated nodes. However, as shown in Fig. 3d and 3g, the amount of storage deployed is strongly dependent on the α value. In particular, on average, when $\alpha = 1.2$ the optimal solution of the ICN planning problem installs 87% less storage than for $\alpha = 0.8$. In other words, for higher alpha values, it is better to deploy more nodes in the network rather than increasing their storage, while the opposite holds for smaller α .

Figures 3e and 3f show the traffic cost component for the Abilene topology, while in 3h and 3i we portray the same metric for Géant. In all of them there is a steep decrease in costs when the budget goes from 1 to $1.5 \max_p$. On the other hand, for larger budgets, very limited improvements are observed, and they are slightly more relevant with $\alpha = 0.8$.

The Zipf popularity exponent has a negligible impact on the cost of the IP network, since it only affects traffic demands for single objects, but not their aggregate value. On the other hand, by comparing the overall cost of the IP network in the two topologies, we can conclude that, on average, Géant leads to a solution 13% more expensive than Abilene. This difference is even more remarkable, especially when considering smaller topologies; for instance, Géant leads to a solution that is 48% more expensive than Netrail, as shown in the full set of plots [18]. In Géant, when $\alpha = 0.8$, the RR heuristic leads to nearly-optimal solutions which are, on average, only 16% more expensive than the optimal counterparts. In line with previous literature [13], [16], ICN allows the operator to reduce his traffic costs remarkably, even when the migration budget is very constrained, saving up to 68% of the overall traffic costs, as we observed in Géant, with $\alpha = 1.2$.

VII. CONCLUSION

In this paper we tackled the *content-aware network planning* problem for the *migration* to an ICN, in a *budget constrained* scenario. In order to derive the optimal strategy that the operator should pursue, we formulated a Mixed Integer Linear Programming model that can be used to jointly identify the node migration strategy, with optimal object placement and request routing. Our proposed optimization model takes into accurate account economic parameters related to: 1) the traffic, 2) the router migration and 3) the caching storage costs. We further complemented our contribution by designing a near-optimal Randomized Rounding (RR) heuristic that scales up to realistic topologies composed of hundreds of nodes.

We discovered that, by migrating only few nodes to ICN, the operator can experience up to a 68% reduction in traffic costs, compared to those of an IP network, as we observed for the Géant topology. On top of that, when the content popularity distribution is very skewed (i.e. $\alpha = 1.2$) the migrated nodes have on average 87% less storage than the one deployed when setting $\alpha = 0.8$. Numerical results show that our proposed RR heuristic can compute close to optimal solutions, and at the same time it can potentially scale up to large network topologies. To the best of our knowledge, this is the first paper that tackles the network planning problem for the migration to an ICN.

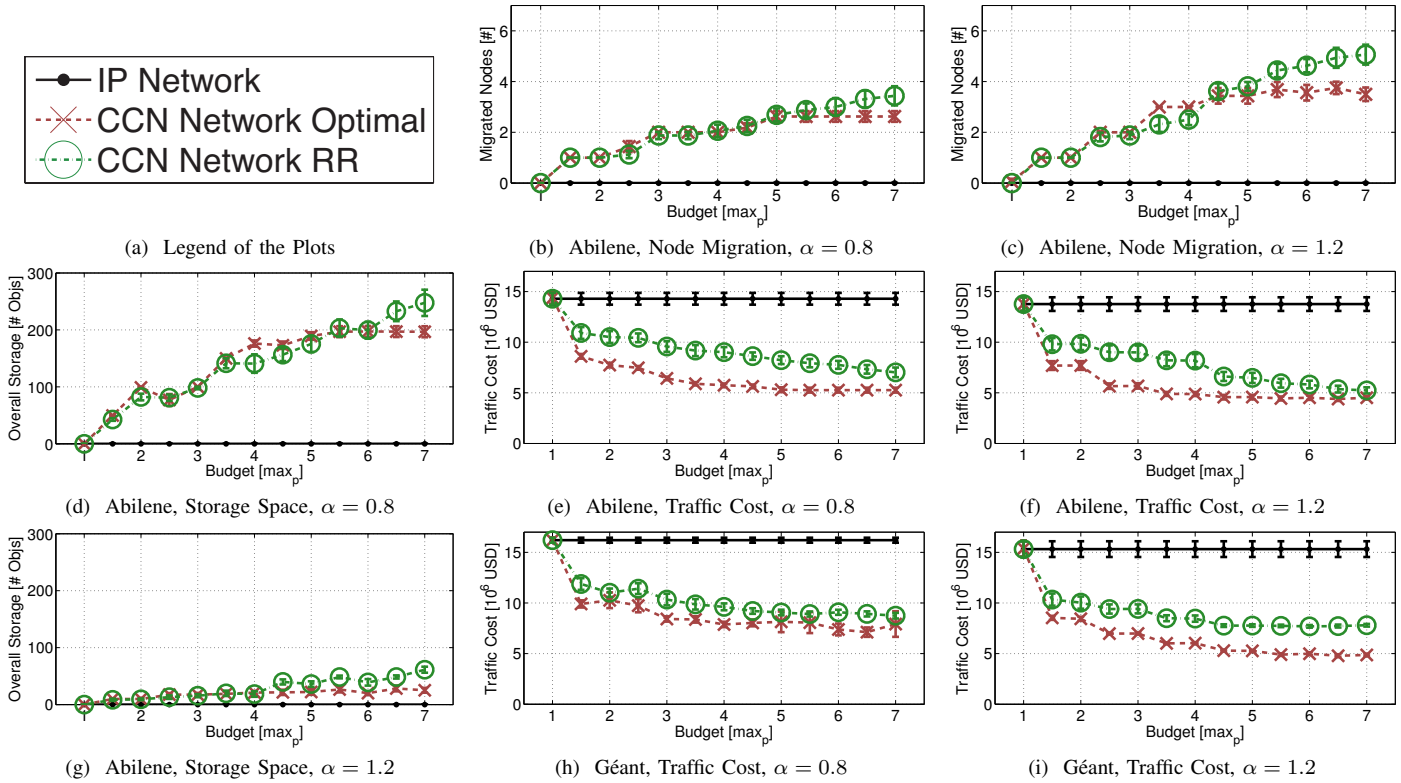


Figure 3. *Effect of the Budget*. Fig. 3a is the common legend of all the plots. Fig. 3b-3g are relative to the Abilene topology, while Fig. 3h,3i refer to Géant. The number of ICN migrated routers is shown in Fig. 3b,3c for different α values, similarly, the overall storage is plotted in Fig. 3d,3g. Lastly, the traffic cost component is depicted in Fig. 3e, 3f for Abilene, and in Fig. 3h,3i the traffic in Géant is available.

REFERENCES

- [1] Sandvine Inc. ULC, “Global internet phenomena report,” 2nd half 2013.
- [2] Cisco, “Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013–2018,” http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf, Last accessed: Apr. 2014.
- [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] G. Carofiglio, G. Morabito, L. Muscariello, I. Solis, and M. Varvello, “From content delivery today to information centric networking,” *Computer Networks*, vol. 57, no. 16, pp. 3116 – 3127, 2013.
- [5] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, “A survey of information-centric networking,” *IEEE Communications Magazine*, vol. 50, no. 7, pp. 26–36, 2012.
- [6] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, “Networking Named Content,” in *Proc. of the 5th Int.l Conf. on Emerging networking experiments and technologies (CoNEXT)*. ACM, Rome, Italy, December 2009, pp. 1–12.
- [7] M. Varvello, D. Perino, and J. Esteban, “Caesar: A content router for high speed forwarding,” in *Proc. of the 2nd Workshop on Information-centric networking (ICN)*. ACM, 2012, pp. 73–78.
- [8] W. So, A. Narayanan, D. Oran, and M. Stapp, “Named Data Networking on a Router: Forwarding at 20Gbps and Beyond,” in *Proc. of the ACM SIGCOMM*, Hong Kong, China, Aug. 2013, pp. 495–496.
- [9] “CCNX Community Meeting (CCNXCon 2013) Keynote,” Palo Alto, CA, USA, Sep. 2013. Speaker: Glenn Edens - <http://www.ccnx.org/events/ccnxcon2013/program/> Last accessed: Apr. 2014.
- [10] S. Arianfar, P. Nikander, and J. Ott, “On content-centric router design and implications,” in *Proc. of the Re-Architecting the Internet Workshop (ReArch)*. ACM, Philadelphia, PA, USA, Nov. 2010, p. 5.
- [11] D. Perino and M. Varvello, “A Reality Check for Content Centric Net-
- [12] P. Krishnan, D. Raz, and Y. Shavitt, “The Cache Location Problem,” *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 568–582.
- [13] Y. Wang, Z. Li, G. Tyson, S. Uhlig, and G. Xie, “Optimal Cache Allocation for Content-Centric Networking,” in *Proc. of 21st IEEE ICNP*, Göttingen, Germany, 7–10 October, 2013, pp. 1–10.
- [14] S. Hasan, S. Gorinsky, C. Dovrolis, and R. K. Sitaraman, “Trade-offs in Optimizing the Cache Deployments of CDNs,” in *To appear in Proc. of the 33rd IEEE Int.l Conf. on Computer Communications (INFOCOM)*, Toronto, Canada, May 2014.
- [15] H. Mihara, D. Imachi, M. Yamamoto, T. Miyamura, and K. Sasayama, “Content Aware Routing: A Content Oriented Traffic Engineering,” in *Proc. of the IEEE Global Communications Conference (GLOBECOM)*, Atlanta, GA, USA, December 2013, pp. 1438–1443.
- [16] M. Mangili, F. Martignon, and A. Capone, “A Comparative Study of Content-Centric and Content-Distribution Networks: Performance and Bounds,” *Proc. of the IEEE Global Communications Conference (GLOBECOM)*, pp. 1425–1431, Atlanta, GA, USA, December 2013.
- [17] O. Klopfenstein, “A Randomized Rounding Heuristic to Reroute Tunnels in MPLS Networks,” in *Proc. of the 5th Int.l Workshop on Design of Reliable Communication Networks (DRCN)*, Ischia, Naples, Italy, Oct. 2005, pp. 461–468.
- [18] “Complete Set of Plots and References to the Considered Topologies,” Available at: <https://www.lri.fr/~mangili/icn-planning.html> Last accessed: Apr. 2014.
- [19] G. Carofiglio, V. Gehlen, and D. Perino, “Experimental Evaluation of Memory Management in Content-Centric Networking,” in *Proc. of IEEE Int.l Conf. on Communications (ICC)*, Kyoto, Japan, June 2011, pp. 1–6.
- [20] C. Fricker, P. Robert, J. Roberts, and N. Sbihi, “Impact of traffic mix on caching performance in a content-centric network,” in *Proc. of IEEE INFOCOM Workshops*, Orlando, FL, USA, March 2012, pp. 310–315.