Performance Evaluation of Video Server Replication in Metro/Access Networks[☆]

Marco Savi[a,*], Roberto Fratini[a], Giacomo Verticale[a], Massimo Tornatore[a]

[a]*Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Piazza Leonardo da Vinci 32, Milano, Italy*

Abstract

Internet traffic is increasingly becoming a media-streaming traffic. Especially, Video-on-Demand (VoD) services are pushing the demand for broadband connectivity to the Internet, and optical fiber technology is being deployed in the access network to keep up with such increasing demand. To provide a more scalable network architecture for video delivery, network operators are currently considering novel metro/access network architectures which can accommodate replicated video servers directly in their infrastructure. When servers for VoD delivery are placed nearer to the end users, part of the traffic can be offloaded from the core segment of the network, and the end users can experience better Quality of Service (QoS). While the deployment of caching systems for traffic offloading has been studied in the core network, no work has already investigated the potential performance gains by replicating the content in the metro/access segment of the network, even closer to the users. In our work, we will compare the performance of video server replication in different metro/access network architectures, i.e. a metro ring architecture and a tree-based architecture, by considering both active and passive technologies. We will evaluate using both simulative and analytical methodologies how content providers could benefit from the deployment of replicas of video servers in terms of blocking probability of the VoD requests.

Keywords: traffic offloading, video-on-demand, content replication, performance evaluation, metro network, access network

---

[☆]A preliminary version of this paper appeared in R. Fratini, M. Savi, G. Verticale and M. Tornatore, "Using Replicated Servers for VoD Traffic Offloading in Integrated Metro/Access Networks", IEEE ICC 2014

[*]Corresponding author. *Email address:* marco.savi@polimi.it, *Tel:* +390223993646

1. Introduction

The constantly increasing adoption of new broadband services such as Media Streaming, File Sharing, Voice over IP (VoIP), and Online Gaming is the main driver of the current Internet traffic explosion. In a recent forecast report [1], Cisco asserts that media streaming traffic is going to become the dominant type of Internet traffic: in particular, Video-on-Demand (VoD) has been recognized as the most prominent media streaming services. Telecom operators are thus striving to sustain the growth pace of such broadband services by deploying new optical access network technologies, such as, e.g., fiber-to-the-home (FTTH) or fiber-to-the-building (FTTB), in order to provide fiber connectivity closer to the end users and higher user access rates [2]. Unfortunately, such increase of user access rates will pose a serious scalability issue to the current telecommunication network infrastructure.

In fact, current telecommunication networks consist of three domains: access, metro/aggregation, core. Core networks provide nationwide/global coverage, and spans long distances, i.e., hundreds to thousands kilometers for each link. Metro networks typically span a metropolitan region and metro switches are usually connected in a ring topology. Access networks connect end users to their service provider, enable end-users to connect to the rest of the network infrastructure, and span a distance of a few kilometers. They are usually connected in a tree topology.

Currently, most video contents are stored in centralized media servers placed in data-centers in the core segment of the network (or directly connected to it) and, without innovative strategies for content delivery, the metro and core segments will soon be flooded by media traffic. Various studies have shown how such segments could constitute a traffic bottleneck that could compromise the performance of VoD services [3][4][5].

A possible strategy to prevent this consists in migrating part of the VoD contents towards the edge of the network (i.e., in the metro and access segments) to offload some of media traffic from the core segment and, possibly, even from the higher layers of aggregation of the metro segment. Such strategy can be achieved by deploying a system of distributed replicated servers in the metro and access networks which ensures fast and decentralized delivery of VoD contents, enabling a more scalable architecture.

Integrated metro/access architectures, where the access and metro infrastructure converge into a single domain (as in the case of Long Reach Passive Optical Networks (LR-PONs) [6]) represent a promising candidate for next-generation cost- and energy-efficient network architectures for content delivery to users [7]. They are typically multi-stage networks spanning up to few hundreds of kilometers. While some studies have already investigated the effect of positioning server replicas in core nodes and/or in access nodes, to the best of our knowledge, a comparative study on the effect of VoD content migration/distribution at various stages of the metro/access network, while considering different metro/access architectures, has never been thoroughly investigated.

In this work, we study the performance of VoD content delivery where Metro Servers (MSs) are placed in the metro/access network. Specifically, we will show how different metro/access network architectures, equipped with different technologies and under different configurations, perform in terms of blocking probability for VoD requests. Especially, we will consider a metro ring and a tree-based architecture, focusing in our evaluation on both active and passive technologies. We will show how exploiting this variety of configurations can lead to performance improvements in terms of blocking probability of VoD requests and how, under simplified assumptions, it is possible to analytically study the traffic bottlenecks occurring in such architectures.

The rest of the paper is organized as follows: in Section 2 we provide an overview of the related work. In Section 3 we describe our representative architectures for server replication and we describe their main differences from a technological perspective. In Section 4 we focus on the modelling of MSs VoD content and traffic. We also provide an overview of the simulator developed for the performance evaluation under our modelling assumptions. In Section 5 we perform a theoretical analysis under simplified assumptions. Then, Section 6 shows and discusses the main results of our performance evaluation. Finally, Section 7 draws the conclusion of this study.

2. Related works

Some existing studies have dealt with content placement and distribution in telecommunication networks. One of the main challenges consists in placing the content servers in different segments of the network while jointly taking into account energy efficiency, core-traffic offloading and latency.

Refs. [8][9][10] are among the first works showing how replicating contents towards the users can help in jointly reducing the latency for content retrieval and improving the energy efficiency of VoD delivery. All these works refer to solutions in which contents are moved closer to the users but still replicated in the core segment of the network. With respect to Refs. [8][9][10], our work mainly focus on the core-traffic offloading (by moving contents even closer to the users in the metro and access networks) and latency aspects while only partially focusing on energy efficiency.

Refs. [11][12][13][14] try to push the content even closer to the users. Ref. [11] introduces the concept of micro Content Delivery Network (micro-CDN), where small servers are placed in some metro nodes, while Ref. [12] analyzes the performance, especially from an energy efficiency perspective, of a centralized VoD distribution system where a single replicated server can be placed in the metro, access or core segment. Ref. [13] deals with performance improvement due to caching in Internet Service Providers (ISPs) networks, while Ref. [14] defines a novel hybrid wireless/optical access network and studies a cost-efficient solution for cache placement over such network. Also our work refers to content delivery in the metro and access segments of the network. However, we focus on a specific performance metric, i.e., the blocking probability of VoD requests, that has not been taken into consideration in such works, and that allows to focus our attention on possible bottlenecks due to infrastructural flaws of the network architectures.

Other works, such as Refs. [4][5][15], pursue peer-to-peer (P2P) solutions for core-traffic offloading by keeping VoD traffic as local as possible in the access network. Our work does not specifically take into consideration P2P solutions. However, in our evaluation, P2P can be seen as a special case, where the contents are pushed to the customer premises. Ref. [4] quantifies the core traffic reduction due to P2P caching when a LR-PON architecture is adopted. Our work, like Ref. [5], does not take into consideration only metro/access passive technologies, but also active technologies. Ref. [15] does not focus on any specific underlying network technology, but defines a strategy for optimal content placement in order to maximize bandwidth utilization. In our work we do not deal with optimal content placement, but we instead offer a tool to dynamically evaluate network performance and bandwidth utilization when specific metro/access network architectures are adopted.

Finally, some innovative content delivery architectures exploiting the novel concepts of Software-Defined Networking (SDN) and Network Functions Vir-

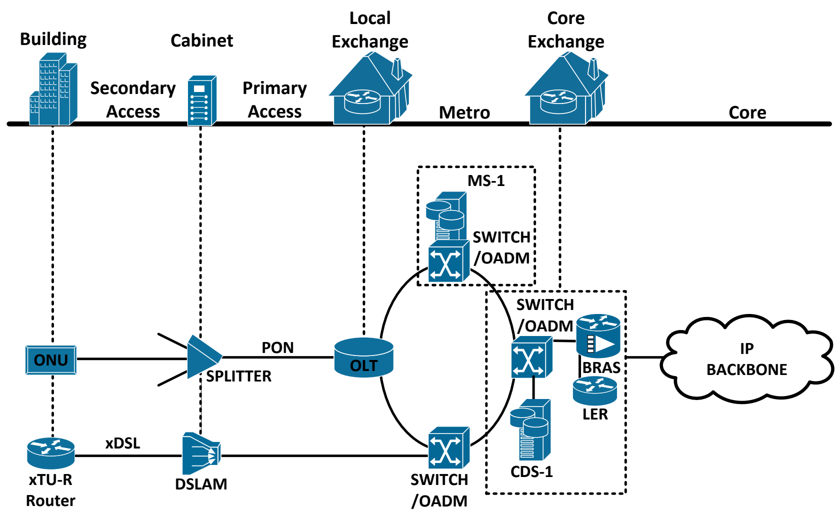

Figure 1: Metro ring network architecture

tualization (NFV) have started being investigated. For example, Refs. [16][17] define OpenCache, a SDN in-network caching service for an efficient content delivery to the users. Instead, Ref. [18] defines a virtualized content cache prototype that can be deployed in micro datacenter close to the users. Even if we do not focus on such aspects, SDN and NFV are very promising concepts to improve caching in the metro/access networks. The proposed solutions of Refs. [16][17] specify a cache coordination scheme that can be easily adopted by the network architectures considered in this paper, as well as the cache virtualization prototype proposed by Ref. [18].

3. Metro/access network architectures for VoD delivery

In this Section, we describe different metro/access network architectures. We first describe an architecture characterized by a metro ring, then we consider a tree-based network architecture and we differentiate between one architecture based mostly on active (e.g., Ethernet) switches and one based mostly on passive optical equipments. For all these solutions, we discuss how replicated servers can be introduced in the network. Note that in this work we refer to generic VoD delivery architectures, which could be adopted by both Subscription VoD (SVoD) and Free VoD (FVoD) content providers. In case of SVoD, some concerns (out of the scope of this paper) should be taken into account, such as the problem of licensing of contents.

3.1. Metro ring network architecture

As shown in Fig. 1, a metro ring network architecture presents a metro ring (e.g., using Gigabit Ethernet) of active (i.e., switches) or passive (i.e., Optical Add-Drop Multiplexers, OADMs) nodes that interconnect different access technologies, such as Digital Subscriber Line (DSL) and/or PON, with the core exchange node, where a Broadband Remote Access Server (BRAS) and/or a Label Edge Router (LER) route and forward traffic to/from the IP backbone network. In this scenario, it is possible to place one or more Metro Servers (as the MS-1 in Fig. 1) in the Central Offices where the switches/OADMs are hosted. Note that in our model, we always assume that a primary content server (referred to as Content Delivery Server-1, CDS-1, and storing all the contents) is placed somewhere in the core network (e.g, at the core exchange)[1].

3.2. Tree-based network architecture

We consider here a tree-based network architecture where MSs can be placed at various stages of the network. In this architecture the metro/access network is characterized by a long-reach branch-tree architecture, as in Fig. 2 and Fig. 3.

3.2.1. Metro Servers and additional mesh link placement

In such scenarios, we introduce the possibility to add m MSs (m-MS configuration) and n mesh links, i.e., links between nodes belonging to the same stage of the tree (n-MESH configuration, see Fig. 2 as an example) to increase the VoD delivery performance. Once an MS is placed in a specific metro node of the network, mesh links placed in strategical positions allow to increase the number of end users that can take advantage of that specific MS. The case with only the CDS-1 and no MS is referred to as 0-*MS configuration.*

3.2.2. Active vs. passive architecture

As shown in Fig. 2 and Fig. 3, the network described above might be implemented using either passive or active optical technology. In the active case,

[1] Note that placing a content server in the core exchange is already a first important step for VoD content migration: in fact, such placement for the CDS-1 solves the core traffic offloading problem. However, in terms of performance evaluation of the metro segment, there is no difference in placing the CDS-1 in the core exchange or somewhere else in the core.

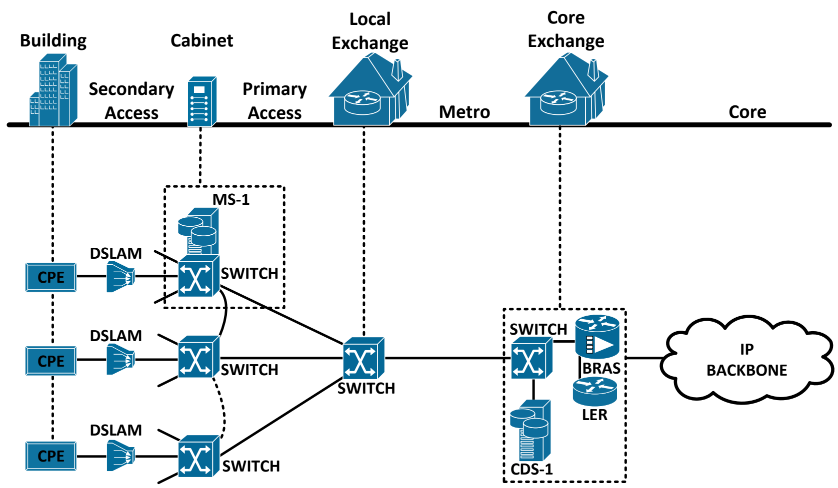

Figure 2: Tree-based *active* network architecture

the nodes of the hybrid meshed/branch tree are Ethernet switches and the addition of DSL Access Mutiplexers (DSLAMs) and Customer Premises Equipments (CPEs) provides broadband access to the end users. In the passive case, an Optical Line Termination (OLT) in the core exchange performs electrical/optical conversion for the downstream traffic, the intermediate metro nodes are power splitters and the Optical Network Units (ONUs) provide broadband connectivity to the end users. There are various arguments that can be used in favor of either solution, and the presence of replicated servers plays a relevant role in choosing the most effective option.

The active solution can more easily incorporate the introduction of an MS in the intermediate stages of this metro/access network. In fact, the active architecture does not require specific upgrades to support the introduction of an MS, while in the passive architecture some significant changes in terms of network architecture are required. Specifically, in the passive case, whenever an MS is introduced in an intermediate metro node, the node is no more passive and an additional active component is needed, e.g. a sub-OLT (see [19] for more details). Such active device works as a hybrid ONU/OLT device, since it acts as an OLT for all the connection requests assigned to it and acts as ONU for all the other traffic which is directed to the sub-OLT's children nodes. For this reason, the sub-OLT enables content delivery from the MS to its children nodes. Note that in this case we refer to the name *passive architecture* also if this is actually a *mostly* passive architecture, in the sense

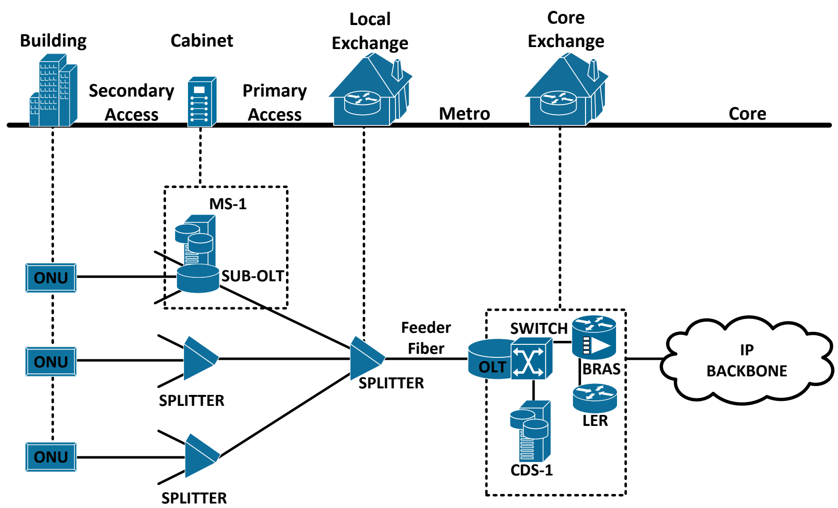

Figure 3: Tree-based *passive* network architecture

that the node hosting the MS cannot be passive, while all the other nodes are.

The addition of mesh links is beneficial only in case an active architecture is deployed. From a technical point of view, it is possible to add mesh links also in the passive architecture by replacing the power splitter with a star coupler [20] in all the passive nodes that are termination of a mesh link. However, because of the broadcasting nature of downstream communication in passive optical networks, the addition of mesh links cannot help in improving network performance in such scenario [21].

If we consider cost and energy efficiency, the passive architecture might represent a more desirable solution than the active one. In fact, in the passive case, all the intermediate nodes which are not equipped with an MS are passive optical nodes, that can lead to savings in terms of energy consumption and other operational expenditures. Conversely, in the active architecture all the intermediate nodes (regardless of the presence of an MS and/or of a mesh link) are equipped with a switch, that is an active device and must be always powered.

We must also consider that an active solution allows more users to be served by a single MS. In fact, switches in an active hybrid meshed/branch tree architecture allow the forwarding of traffic from the MS also in the upstream direction, therefore all the end users of the metro/access network can be reached by an MS placed anywhere in the metro/access network. This

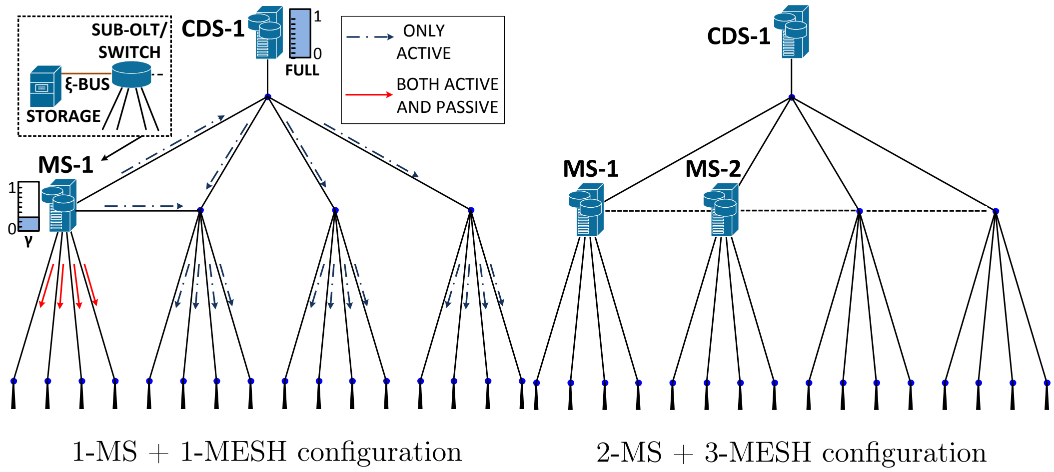

Figure 4: Example of a 64-NTs (64-Network Terminations) tree-based network

is shown in Fig. 4, where red and dotted blue arrows depict the possible directions of flows from the MS to end users in the passive and in the active case. Note that such difference in terms of flows and reachability of users between an active and a passive architecture does not exist in the metro ring network as described in Section 3.1. As stated before, the metro ring can be deployed both using active technologies (i.e., switches) or passive technologies (i.e., OADMs), but in both cases flows in the ring move in the same possible directions (clockwise and counter-clockwise).

Finally, it is important to note that Fig. 1, Fig. 2 and Fig. 3 are just examples of reference architectures, where the MSs are placed in central offices (Fig. 1) and in street cabinets (Fig. 2 and Fig. 3). Clearly, depending on the geographical and hierarchical structure of the actual network, the topology could vary and the MSs could be placed in different physical locations.

3.3. Preliminary cost analysis

Although a detailed cost analysis is outside the scope of this paper, we performed some preliminary cost considerations. The deployment of an m-MS configuration leads to some increased capital expenditure (CAPEX) that cannot be avoided. According to [22], a *volume server* costs less than 20000$ and in average consumes 222 W. We assume a cost of 5000 $ for each MS, and thus a CAPEX of $m \cdot 5000$ $ for am m-MS configuration. Concerning operational expenditures (OPEX), we assume that the MSs are always on and we consider an average price per kilowatt-hour of 0.12 $/kWh [23]. This leads to a cost of $\simeq 235$ \$/year for each MS, and thus an OPEX of $m \cdot 235$ \$/year for a m-MS configuration. Different strategies to reduce OPEX can be investigated. For example, in Ref. [24] we show how MSs can be switched on/off according to traffic load variations through the day in order to reduce the overall energy consumption of the VoD content delivery service.

In the next Section, we will define a practical modelling of the VoD service, which will be our reference model for the performance evaluation.

4. Simulative modelling of server replication

The aim of this paper is to evaluate the performance of the VoD service when it is deployed in different metro/access network architectures. Specifically, given:

1. A metro/access network architecture (metro ring, active tree-based or passive tree-based);
2. A specific configuration in terms of placed MSs (i.e., how many and where they are placed);
3. A statistical modelling of the VoD content;
4. A statistical modelling of the VoD traffic;
5. A modelling of the MSs in terms of *content replicability* and *capacity*;

we want to evaluate the *blocking probabilty* of the VoD content requests, defined as the number of VoD content requests that cannot be accommodate due to lack of resources with respect to the overall number of VoD content requests generated by the VoD users in a specific time-interval (for a formal definition of blocking probability see Section 4.3).

After having described in Section 3 the metro/access network architectures, in this Section we will describe how we model the VoD content, the VoD traffic and the MSs. We will then describe the simulator we developed for the simulative performance evaluation based on our modelling.

4.1. VoD content and traffic modelling

Different models to represent VoD contents and traffic can be used. In this paper we use results from [25], where a long-term trace of a real VoD streaming traffic is collected and statistically analyzed. We model the following properties of a VoD system:

- *Connection requests:* According to [25], the arrival of video requests can be modelled as a time-varying Poisson process. This process is characterized by an arrival rate λ_i, where the index i indicates the hour of the day. This way we consider a non-stationary process where prime time traffic (defined as the peak traffic hour) is modelled by generating connection requests arrivals at a rate roughly 1.5 times higher than the average arrival rate.
- *Content duration:* Content duration is largely determined by the nature of the content stored in the content servers, i.e. short/long clips [25], [26]. Thus, the content duration distribution results in a multimodal distribution having multiple peaks. In our work a finite mixture-gaussian as defined in [25] is used in order to capture different peaks of the duration distribution for short/long clips. The probability density function of a generic mixed-gaussian distribution is the following:

$$g(x) = \sum_{i=1}^{n} w_i \cdot \frac{1}{\sigma_i \sqrt{2\pi}} \cdot e^{-\frac{(x-\mu_i)^2}{2\sigma_i^2}}$$

  where n is the number of gaussian distributions to be mixed, each one capturing a different peak of duration, the i-th gaussian distribution has mean μ_i and variance σ_i^2, and w_i is the weight associated to the i-th gaussian distribution. In this paper we consider two different pdfs for content distribution, i.e., $g_{short}(x)$ and $g_{long}(x)$, characterized by a different mix of weighted gaussian functions [25]. We assume that half of the contents in the catalogue are short clips and half long clips.
- *Content popularity:* Video popularity is statistically modelled through the Zipf distribution [27], which is often used for content popularity of large VoD media systems. In general, if M is the number of media contents provided by the VoD content provider through the CDS-1, 1 is the rank of the most popular content and M the rank of the least popular, the probability that a generic content with rank $1 \leq m \leq M$ is requested by an user is defined as:

$$h(m) = \frac{K}{m^{\phi}}$$

  where K is a normalization constant and ϕ is a distribution parameter. In this work, we consider $\phi = 1.5$.
- *Encoding bitrate:* A set of common used streaming bitrates is defined in order to model a discrete distribution of constant bitrates, as reported

Quality	Video	Audio	Bitrate
SD	640x360 px	Stereo @ 44kHz	1.3 Mbit/s
HQ	960x540 px	Stereo @ 44kHz	2.8 Mbit/s
Mobile	960x640 px	Stereo @ 44kHz	3.3 Mbit/s
HD	1280x720 px	Stereo @ 44kHz	4.9 Mbit/s

Table 1: Bitrates for different video streaming settings

in Table 1. We suppose that the bitrate chosen by a user when a content is requested is uniformly distributed among these bitrates.

4.2. MSs modelling

We also model two important features of the MSs, i.e., the storage content replicability that an MS can offer and the bus capacity of the MS:

- *Content replicability:* We define the ratio of replicability γ (with $0 \leq \gamma \leq 1$) as the ratio of the CDS-1 contents that are replicated in an MS (see Fig. 4). The most popular contents, according to the Zipf distribution, are replicated first. Note that if $\gamma = 0$ no content is replicated in the MSs, while if $\gamma = 1$ all the contents are replicated regardless of their popularity, and that usually $\gamma \ll 1$ in a realistic scenario, since the storage capability of MSs is usually much lower than the storage capability of CDS-1.
- *Capacity:* Each MS is connected to the transmission device by a generic bus. Note that the bus could represent a limitation since its capacity might not meet the demand which the MS must satisfy. In order to investigate such limitation, we introduce the ratio of capacity ξ, which represents the capacity of the bus in relation to the capacity of a network link:

  $$\xi = \frac{C_{bus}}{C_{link}}$$

  The ξ-bus is depicted in Fig. 4. We will see how the bus link may result in a bottleneck for the MS, especially if its capacity is not correctly dimensioned. Note that $0 \leq \xi < \infty$. Considering C_{link} as a fixed value, this means that if $\xi = 0$ the MS is disconnected by the network, while if $\xi = \infty$ the ξ-bus cannot be a bottleneck for the MS.

A pictorial description of the capacity and replicability constraints is reported in Fig. 4, where a 64-NT network model is depicted. For a detailed

# ID	Duration	Bitrate	Popularity	Hosting MS
i	*mixed-gaussian*	*uniform extraction*	*Zipf*	[...]

(a) Generic modelling for content i

# ID	Duration (s)	Bit-rate (Mbit/s)	Popularity	Hosting MS
1	1200	2.8	0.38	1;3
2	300	1.3	0.13	1;3
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
n	2200	4.9	$3.9 \cdot 10^{-4}$	1

(b) Example of content generation

Table 2: Generation of content

description of the network topologies used in our simulations please refer to Section 6.

4.3. Overview of the simulator

A discrete-events simulator for performance evaluation has been developed using Matlab. The simulator evaluates performance in terms of blocking probability of VoD requests, considering that a connection between a content server and an user has to be established for each VoD request.

First of all, a fixed number of contents are generated according to the models shown in Section 4.1. We consider a catalogue of $N_c = 10000$ video contents. Table 2 shows an example of content generation together with the model used for each content property. Note that each content is also described by the *Hosting MS* parameter, i.e., the servers storing that specific content. Only $\gamma \cdot N_c$ contents are replicated in the MSs.

Then, every connection request and release is represented by an event scheduled in an event list. Table 3 shows an example for the event list. Each event is characterized by some parameters: *# Flow* is a progressive ID number that characterizes each VoD streaming flow, *Event type* indicates whether the event is a connection request (i.e., *start*) or a connection release (i.e., *end*), t_{event} indicates the timestamp when the event occurs. Clearly, each VoD streaming flow is characterized by a *start* and an *end* event. Especially, the time interval between two consecutive *start* events depends on the inter-arrival time τ between connection requests, while the time interval between a *start* and *end* event for the same streaming flow (i.e., with the same #

# Flow	Event type	t_{event}	Content	Source	NT	Path
1	*start*	2.1 s	4	1	26	[1 2 3 7 26]
2	*start*	3.2 s	10	3	23	[3 7 23]
2	*end*	65.2 s	10	3	23	[3 7 23]
1	*end*	102.4 s	4	1	26	[1 2 3 7 26]
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Table 3: Structure and example of the event list

Flow) depends on the duration of the content that must be served. *Content* indicates the identificator (ID) of the requested content (as reported in Table 2b), *Source* indicates the MS assigned to the streaming flow, *NT* is the ID of the Network Termination requiring the content (we assume that a termination can serve multiple users at the same time) and *Path* reports the path followed by the streaming flow, determined by the *nearest routing algorithm*. The nearest routing algorithm checks the paths between the user requesting the content and the MSs/CDS-1 that store the requested content starting by the shortest in terms of number of hops. If the links of the path between the selected MS and the user can provide enough bandwidth to serve the request (i.e, all the links are available), the MS is marked as source for the current request (i.e., the MS is available) and the bandwidth is allocated. Otherwise, another path between an MS and the user, if any, is considered. If no path from any MS can allocate resources for the current request, the request is forwarded to the CDS-1. If the CDS-1 is unavailable, i.e., no path to the CDS-1 can allocate resources for the request, the request is discarded (i.e., *blocked*).

Note that the simulations stop only when a desired level of statistical confidence has been reached. In this paper, all the plotted values have a 95% confidence interval not larger than 5% of the plotted value.

5. Analytical modelling of server replication under simplified assumptions

In this Section we provide a theoretical evaluation of the blocking probability of connection requests for different network configurations. With respect to the content and traffic modelling parameters shown in Section 4, our theoretical analysis is performed under simplified assumption. This way, we can manage closed-form formulas to calculate the blocking probability of connection requests. Specifically, we assume that:

1. Each content has the same popularity (i.e., every content has the same probability to be requested by a user).
2. Each content is encoded by a single bitrate R_b, and the bitrate is the same for all the stored content.
3. The Poisson distribution used to model the probability distribution function of video requests does not vary through the hours of the day.
4. Content duration in modelled by an exponential distribution whose average value is fixed at a value of $1/\mu$.

These simplified assumptions allow us to model the content delivery service exploiting queuing theory for homogeneous traffic, where we consider as traffic the user VoD connection requests that must be accommodate by the VoD system. We can model each link as a $M/M/s/0$ queuing system, where for the number of servers m we can assume a value of $s = C_{link}/R_b$. Note also that the queues have no buffer, since if all the servers are busy when a request arrives, such request is dropped. In this context, the *blocking probability* of VoD connection requests with respect to a single link is defined as the probability that all the servers of the system are busy while a new request is mapped to that link, i.e., the classic Erlang-B formula:

$$P_b(A_{link}) = \frac{\frac{A_{link}^s}{s!}}{\sum_{k=0}^{s} \frac{A_{link}^k}{k!}} = E_{B,s}(A_{link}) \quad (1)$$

where A_{link} is the average offered load to the link. More generally, the blocking probability of the whole VoD system can be analytically obtained by investigating the *bottleneck* of the architecture, i.e., the link (or the set of links) whose block is the highest in the whole system. Our theoretical formulations build upon the classical *overflow traffic theory* [28].

In the remainder of this Section we derive the analytical formulas for blocking probabilities considering different configurations in terms of MSs (m-MS) for both the tree-based network architecture (considering both the passive (PAS) and active (ACT) cases) and the metro ring architecture (RING). Note that for the tree-based active architecture we do not consider mesh link addition. This means that we always consider a tree topology for both the tree-based passive and active architectures. However, our theoretical analysis can be easily extended also to that more generic case of hybrid meshed/tree topologies. Note also that although the simplified assumptions of this Section could lead to distorted results in terms of blocking probability, our analytical evaluation allows to benchmark the simulator quality and gives the possibility to achieve generic guidelines and trends concerning the evaluation of the bottleneck links in the system for different network architectures and configurations.

5.1. m-MS unconstrained configuration ($\gamma = 1$, $\xi = \infty$)

We start by assuming $\gamma = 1$ and $\xi = \infty$, i.e., the MSs replicability is not constrained (all the content can be replicated by the MSs) and the MSs bus capacity is not a bottleneck for the system.

5.1.1. Tree-based passive architecture (m-MS,PAS)

Changing the number of MSs leads to a traffic redistribution over the links. The main effect of traffic redistribution is to offload traffic from the CDS-1 and the feeder fiber to the MSs. Note that, when no MS is deployed in the network (0-MS configuration), all the requests must be handled by the CDS-1. So, in a generic tree topology, the bottleneck is the feeder link. For this reason, in the 0-MS configuration we can express the blocking probability in the following way:

$$P_b(A_o) = E_{B,s}\left(A_o\right) \tag{2}$$

where $A_o = \lambda/\mu$ is the overall traffic load offered to the network.

When one or more MSs are introduced in the architecture and thus a m-MS,PAS architecture is considered, we can express the blocking probability in the following way:

$$P_b(A_o) = \frac{r^{l-i} - m}{r^{l-i}} E_{B,s}\left(\frac{r^{l-i} - m}{r^{l-i}} A_o\right) \tag{3}$$

where r is the branching factor (outdegree) of the tree, l is the number of the stages of the tree, i is the number of the stage where the MS are placed, starting from the deeper level (i.e., $i = 0$ means that the MSs are placed at the customer premises) and m is the number of MSs placed. We assume that the m MSs can be placed in any position but only in one level of the tree. Inequalities $m \leq r^{l-i}$ and $0 \leq i < l$ always hold. For example, the right-hand side of Fig. 4 depicts a tree topology where $r = 4$, $l = 3$, $i = 2$ and $m = 2$.

Eq. 3 shows that placing one or more MSs at some stage of the tree allows to reduce the blocking probability of the video requests. In fact, the bottleneck link, in the passive configuration, is always the feeder link, since no upstream traffic is allowed, as depicted in the left-hand side of Fig. 5.

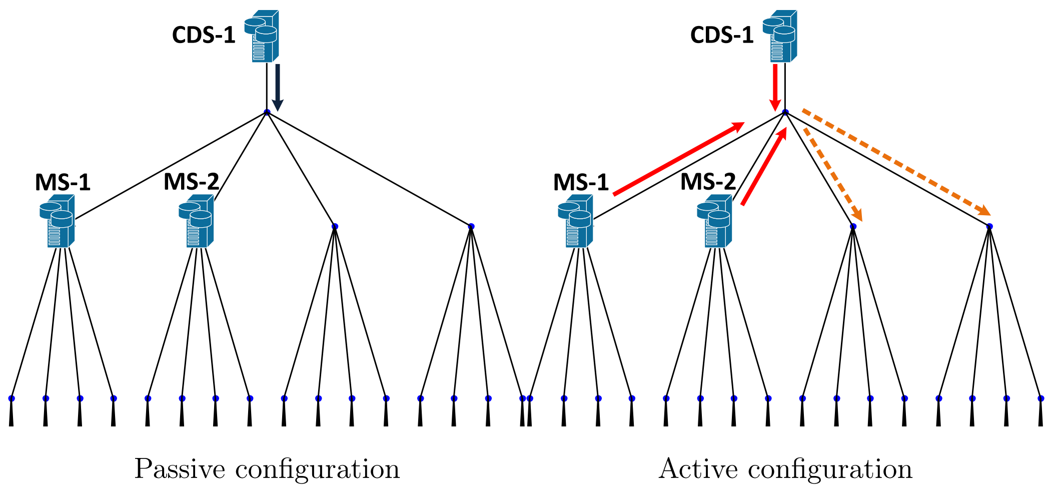

Figure 5: Possible bottlenecks occurring in the 2-MS,PAS and 2-MS,ACT configurations

Placing the MSs simply allows to reduce the video traffic on the feeder link of the quantity $\frac{r^{l-i}-m}{r^{l-i}} \leq 1$.

5.1.2. Tree-based active architecture (m-MS,ACT)

Considering the active architecture, theoretical formulations for the blocking probability are different. In fact, as pointed out in Section 3, the active switches in such architecture enables video request flows in the upstream direction (see the left-hand side of Fig. 4). For this reason, the bottleneck of the architecture is not the feeder link, as in the case of the passive architecture. In this case, we can express the blocking probability in the following way:

$$P_b(A_o) = \max\left\{ \frac{r^{l-i}-m}{r^{l-i}} E_{B,(m+1)s}\left(\frac{r^{l-i}-m}{r^{l-i}} A_o\right); \frac{(r - \lceil \frac{m}{r^{l-i-1}} \rceil) \bmod r}{r} E_{B,s}\left(\frac{1}{r} A_o\right) \right\} \quad (4)$$

Depending on the number m of MSs and on their position, two different bottlenecks can occur. The first bottleneck are the feeder link and the upstream links directly connected to the placed MSs considered as a single queue, that can be modeled as a $M/M/(m+1)s/0$ queue. The feeder link and upstream links considered as a single queue are a bottleneck for the system only when a small number of MSs is deployed. Otherwise, the bottleneck can be detected

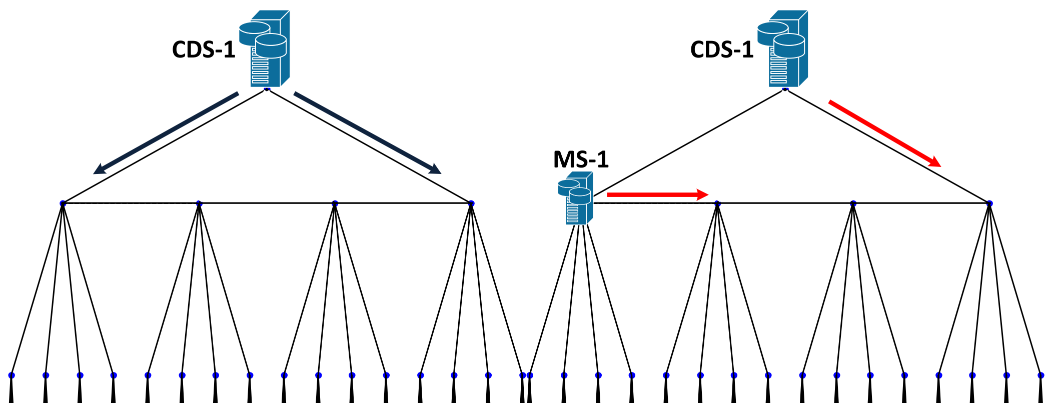

Figure 6: 1-MS, RING and 2-MS, RING ring topologies

in the downstream links connected to the switch placed in the root of the tree and not connected to an MS in the lower level. With respect to the passive architecture, where our formulation is valid for each position for the m MSs at a specific stage, in our formulation of Eq. 4, we assume that the MSs are always placed from left to right in the stage i of the tree. As an example, the right-hand side of Fig. 5 depicts the possible bottlenecks that can occur considering such configuration: the red solid lines and the orange dotted lines depict respectively the first and the second possible bottlenecks.

5.1.3. Metro ring architecture (m-MS,RING)

With respect to the active and passive tree-based architectures, in the metro ring architecture a feeder link does not exist. We can express the blocking probability in the following way:

$$P_b(A_o) = \begin{cases} \frac{1}{2}E_{B,s}(\frac{1}{2}A_o) + \frac{1}{2}E_{B,s}(\frac{1}{2}A_o) = E_{B,s}(\frac{1}{2}A_o) & \text{if } m = 0 \\ \frac{q-m-1}{q-1}E_{B,2s}(\frac{q-m-1}{q-1}A_o) & \text{if } m \geq 1 \end{cases} \quad (5)$$

where q is the number of nodes of the ring. For example, the ring in Fig. 6 has $q = 5$. Note that the expression of the 0-MS configuration is different with respect to the 0-MS configuration of the tree-based active and passive architectures. Looking at Fig. 6, it is easy to see how, in this case, the bottlenecks are the two downstream links connected to the CDS-1, each one allocating half of the global traffic generated by the CDS-1. When one or more MSs are added to the metro ring topology, the bottleneck location changes. In this case, the bottlenecks are the links connecting the rightmost MS to the first metro node without MS and the right link connected to the CDS-1 considered as a single queue. Note that in Eq. 5 and, more generally, in all the evaluations reported in this paper, we assume that the MSs are always placed increasingly in a counter-clockwise direction. Even though this assumption allows to parametrize Eq. 5, placing the MSs in such way is suboptimal. In fact, blocking probability can be further reduced by fairly placing the MSs on the ring, i.e., by minimizing the average distance of the users from their two closest MSs. This way, the bottleneck links experience a lower load. A deeper discussion on optimal placement of MSs in a metro ring is left for future works.

5.2. m-MS capacity constrained configuration ($\gamma = 1$, $0 \leq \xi < \infty$)

Here we consider the capacity constraint as introduced in Section 4.2. A capacity constrained configuration must take into account the ξ-bus as a new possible bottleneck.

5.2.1. Tree-based passive architecture (m-MS,PAS)

We can express the blocking probability in the following way:

$$P_b(A_o, \xi) = \frac{r^{l-i} - m}{r^{l-i}} E_{B,s} \left(\frac{r^{l-i} - m}{r^{l-i}} A_o \right) + \frac{m}{r^{l-i}} E_{1,\xi s} \left(\frac{1}{r^{l-i}} A_o \right) \tag{6}$$

So, the blocking probability of the system is the sum of the blocking probability of the feeder link, like in Eq. 3, and the blocking probability of all the ξ-bus of the MSs placed in the architecture. Note that the ξ-buses of the MSs behave according to a $M/M/\xi s/0$ queue. It is worth noting that $\lim_{\xi \to \infty} E_{1,\xi \cdot s}(\frac{1}{r^{l-i}} A_o) = 0$. This means that for high values of ξ the ξ-bus is not a bottleneck of the system.

5.2.2. Tree-based active architecture (m-MS,ACT)

For the active configuration, the formulation is similar to Eq. 4. We can write:

$$P_b(A_o, \xi) = \max \left\{ E_{B,s(\xi m+1)}(A_o); \frac{r^{l-i} - m}{r^{l-i}} E_{B,(m+1)s} \left(\frac{r^{l-i} - m}{r^{l-i}} A_o \right); \right.$$
$$\left. \frac{(r - \lceil \frac{m}{r^{l-i-1}} \rceil) \bmod r}{r} E_{B,s} \left(\frac{1}{r} A_o \right) \right\} \tag{7}$$

The first term is dominant when the bottleneck is represented by the feeder link and all the ξ-buses. Since a single MS, in the active configuration, can potentially serve all the users of the metro/access network, in this case the system can be modelled as a $M/M/s(\xi m+1)/0$ queue. The second and the third term, that are the same as in Eq. 4, assure that if the ξ-buses are not a bottleneck for the system (i.e., the value of ξ is big enough), the bottleneck can be identified in the feeder link or in the downstream links aggregating upstream traffic directed to the CDS-1 or to the MSs, as in Eq. 4 and Fig. 5.

5.2.3. Metro ring architecture (m-MS,RING)

The blocking probability can be expressed in the following way:

$$P_b(A_o,\xi) = \max\left\{E_{B,s(\xi m+1)}\left(\frac{1}{2}A_o\right); \frac{q-m-1}{q-1}E_{B,2s}\left(\frac{q-m-1}{q-1}A_o\right)\right\} \quad (8)$$

The first term is dominant when the bottleneck is represented by the two downstream links connected to the CDS-1 and all the ξ-buses of the MSs deployed in the ring network. In this case, the system can be modelled as a $M/M/s(\xi m+1)/0$ queue. The second term, which is the same term as the term in Eq. 5 for $m \geq 1$, is dominant when the ξ-buses do not influence the performance of the system.

5.3. m-MS replicability constrained configuration ($0 \leq \gamma \leq 1$, $\xi = \infty$)

Here we consider in our formulations the replicability constraint, as introduced in Section 4.2. A value of $\gamma < 1$ increases the amount of requests that need to be served by the CDS-1, since only the CDS-1 can meet the video requests for the contents that are not stored in the MSs.

5.3.1. Tree-based passive architecture (m-MS,PAS)

As shown in Section 5.1.1, in the passive architecture the bottleneck link is always the feeder link. For this reason, we can express the blocking probability in the following way:

$$P_b(A_o,\gamma) = \left(\frac{r^{l-i}-m}{r^{l-i}} + \frac{m}{r^{l-i}}(1-\gamma)\right)E_{B,s}\left(\left(\frac{r^{l-i}-m}{r^{l-i}} + \frac{m}{r^{l-i}}(1-\gamma)\right)A_o\right) \quad (9)$$

With respect to Eq. 3, the replicability constrained configuration for the passive architectures adds the $\frac{m}{r^{l-i}}(1-\gamma)$ term. This term indicates that a fraction $1-\gamma$ of the connection requests coming from users that could potentially be served by an MS cannot be accommodate by the MS but must be forwarded to the CDS-1, because the MS does not store the requested content.

5.3.2. Tree-based active architecture (m-MS,ACT)

The blocking probability can be expressed in the following way:

$$P_b(A_o,\gamma) = \max\left\{(1-\gamma)E_{B,s}\left((1-\gamma)A_o + \frac{r^{l-i}-m}{r^{l-i}}\gamma E_{B,s}\left(\frac{r^{l-i}-m}{r^{l-i}}\gamma A_o\right)\right);\right.$$
$$\left.\frac{(r-\lceil\frac{m}{r^{l-i-1}}\rceil) \bmod r}{r}E_{B,s}\left(\frac{1}{r}A_o\right) + \frac{\lceil\frac{m}{r^{l-i-1}}\rceil \bmod r}{r}(1-\gamma)E_{B,s}\left(\frac{1}{r}(1-\gamma)A_o\right)\right\} \tag{10}$$

This formulation has two terms. When the first term dominates, the bottleneck of the system is the feeder link. The CDS-1, in this case, has to accommodate the $1-\gamma$ ratio of traffic that cannot be accommodate by the MSs (because they do not store the requested content) and the overflow traffic related to the requests that are not served by the MSs. When the second term dominates, the bottlenecks are the downstream links, like in Eq. 4, but considering also the traffic on the downstream links directly connected to the MSs that cannot be accommodate by the MSs themselves.

5.3.3. Metro ring architecture (m-MS,RING)

The blocking probability can be expressed in the following way:

$$P_b(A_o,\gamma) = \max\left\{(1-\gamma)E_{B,s}\left(\frac{1}{2}A_o\left((1-\gamma)+\frac{1}{2^m}\gamma A_o\right)\right);\right.$$
$$\left.\frac{q-m-1}{q-1}E_{B,2s}\left(\frac{q-m-1}{q-1}A_o\right)\right\} \tag{11}$$

The formulation has two terms. When the first term dominates, the bottlenecks of the system are the downstream links connected to the CDS-1. In this case, the CDS-1 must serve a ratio $1-\gamma$ of the traffic that cannot be served by the MSs (because they do not store the requested content) and the ratio γ of traffic that is served by the CDS-1 independently by the placement of the MSs. Instead, the second term dominates when the bottlenecks become the links connecting the rightmost MS to the first metro node without MS and the right link connected to the CDS-1 considered as a single queue.

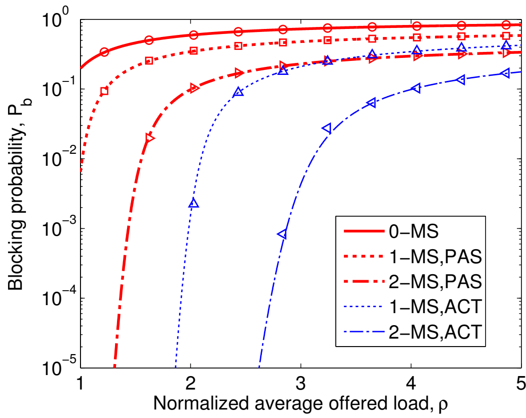

Tree-based architecture

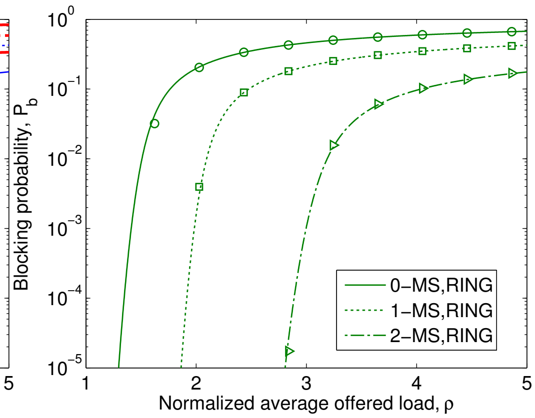

Metro ring architecture

Figure 7: Blocking probability for increasing load in different m-MS unconstrained configurations ($m = 1, 2,\ \gamma = 1,\ \xi = \infty$). *Markers:* simulations, *lines:* theory

6. Performance evaluation

We first show simulative and analytical results in the simplified scenario as defined in Section 5. Then, we show simulation results in the realistic scenario, as modelled in Section 4. In all our simulations we consider the topologies shown in Figs. 4, 5 and 6. Figures 4 and 5 show two generic tree-based network architectures, with a per stage branching factor of 1:4. In this case, the overall network consists of three stages leading to a total branching factor of 1:64 (64-NTs network model). Figure 6 shows a generic metro ring architecture. Sub-trees are connected to the ring nodes, and also in this case the network connects 64 NTs. In both the tree-based and ring architectures, we assume that the capacity for each link progressively increases while considering upper layers of the topology. In the topologies shown in Fig. 4 and 6 we consider that the end users are directly connected to links of 10 Mbit/s capacity, 10 Mbit/s links are aggregated into 100 Mbit/s links and upper layers links (both in the tree and in the metro ring) have a capacity of 1 Gbit/s.

6.1. Analytical vs. Simulation results in the simplified scenario

In the simplified scenario, we consider a bitrate $R_b = 3.075$ Mbit/s for all the stored content. This value is the average bitrate with respect to the four bitrates in Table 1. We also consider an exponential distribution for the content duration whose average value $1/\mu = 2200$ s. This value is the average duration of a content when the content duration is distributed as shown in Section 4.1. Figure 7 shows the blocking probability P_b of VoD connection requests as a function of the normalized average offered load ρ, by assuming that traffic is uniformly distributed among the NTs, in case of $m = 0, 1, 2$ for the tree-based passive, tree-based active and metro ring architectures. We consider an unconstrained configuration ($\gamma = 1$, $\xi = \infty$). In the graphs, the markers represent the simulation points while the lines show the theoretical curves. The theoretical analysis closely approaches the simulation results. The introduction of one or more MSs in the architecture can consistently reduce the blocking probability with respect to the 0-MS configuration in all the architectures. The tree-based passive architecture performs worse than the tree-based active and the metro ring architectures, that in turn behave in a similar way, except for the 0-MS and 0-MS,RING curves, where the ring topology performs better because the bottleneck is not a single link but two links. Considering the tree-based architectures, the 0-MS configuration clearly performs in the same way for both the active and passive cases.

The left-hand side of Fig. 8 shows the blocking probability P_b for the connection requests when the average traffic load is fixed at the value $\rho = 3$ and the ratio of capacity ξ varies in the range $0 \leq \xi < \infty$ in case of an m-MS configuration, considering $m = 0, 1, 2$. In the definition of ξ, we consider $C_{link} = 1$ Gbit/s. Also in this case, we can see how simulation results closely match the analytical curves. We can notice how the blocking probability decreases as ξ increases until a knee point is reached. For example, in the active architecture, we can notice how the knee point appears for a value around $\xi = 0.7$ in case of the 1-MS,ACT configuration, and around $\xi = 0.8$ while considering the 2-MS,ACT configuration.

The right-hand side of Fig. 8 shows the blocking probability P_b for the connection requests when the ratio of replicability γ varies in the range $0 \leq \gamma \leq 1$ in case of a m-MS configuration, considering $m = 0, 1, 2$. We can see how simulation results perfectly match the theoretical analysis and how an increase in the parameter γ leads to a linear decrease in the blocking probability for both the tree-based passive configurations, i.e., the 1-MS,PAS and the 2-MS,PAS configurations. Considering instead the active configurations, i.e., the 1-MS,ACT and the 2-MS,ACT configurations, we can notice a knee point around the value $\gamma = 0.6$. The same occurs for the metro ring configuration, with the knee point around $\gamma = 0.4$ for the 1-MS,RING configuration and $\gamma = 0.5$ for the 2-MS,RING configuration.

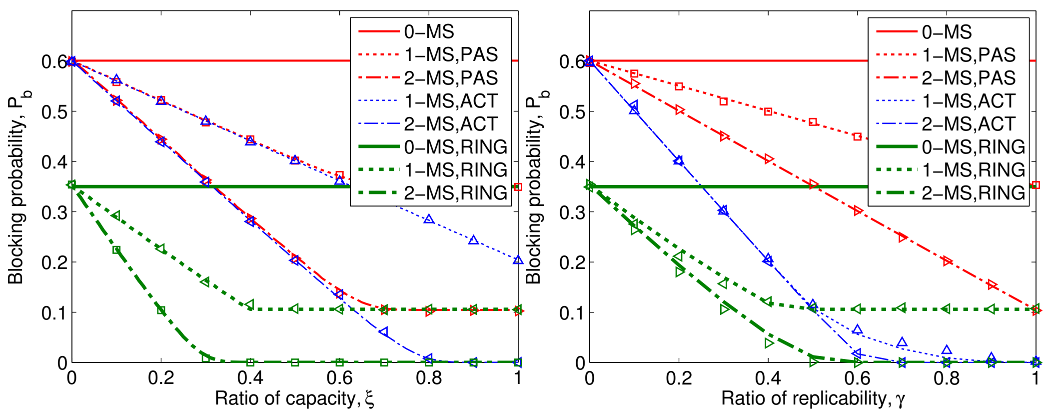

$\gamma = 1$, $0 \leq \xi < \infty$, $\rho = 3$ $\qquad$ $0 \leq \gamma \leq 1$, $\xi = \infty$, $\rho = 3$

Figure 8: Blocking probability in m-MS capacity and replicability constrained configurations in the ideal scenario. *Markers:* simulations, *lines:* theory

The appearance of knee points in Fig. 8 for both the capacity and replicability constrained configurations is justified by the change, for specific values of γ and ξ, in the set of bottleneck links, as described in Sections 5.2 and 5.3.

In general, an average traffic load of $\rho = 3$ leads to blocking probabilities that are extremely high for the performance of a real VoD system, but allows to better visualize the knee for the different configurations. Analogous considerations can be done for smaller values of ρ.

6.2. Simulation results in the realistic scenario

In this Section, we relax the assumptions of the previous Section and we consider a more sophisticated VoD system modelled as discussed in Section 4. First of all, we consider the 1-MS configuration. This configuration is characterized by a single MS placed as shown in the left-hand side of Fig. 4 for the tree-based active/passive architectures and in the right-hand side of Fig. 6 for the metro ring architecture. In case of a tree-based active architecture, the topology can be further modified by adding up to 3 mesh links (1 MS + 1/2/3-MESH,ACT configurations). The blocking probability P_b related to such network configurations when a single MS is deployed is shown in the left-hand side of Fig. 9. We can see how in the 1-MS,PAS case P_b decreases about an order of magnitude with respect to the 0-MS case. In addition to this, we might observe how P_b significantly decreases when we consider

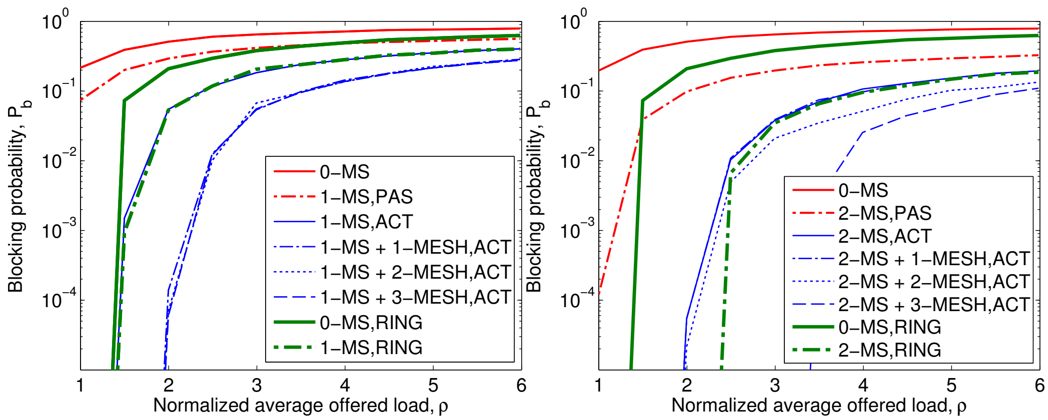

$\xi = \infty$, $\gamma = 1$, 1-MS $\quad\quad$ $\xi = \infty$, $\gamma = 1$, 2-MS

Figure 9: Blocking probability for increasing load in different 1-MS and 2-MS configurations in the realistic scenario

the tree-based active architecture with respect to the same configuration in the passive architecture, since part of the traffic on the feeder link is further offloaded to the MS. A further significant decrease of P_b is achieved in the 1-MS + 1-MESH,ACT configuration with respect to the 1-MS,ACT case. In fact, in this configuration the MS-1 is able to send video traffic either by using the upstream direction or by using the mesh link. Finally, minimal improvements are achieved while increasing the number of mesh links for the active case. Considering instead the metro ring architecture, we can see how the 0-MS,RING topology leads to better results than the 1-MS,PAS configuration but worse result than all the considered configurations for the tree-based active architecture. Then, the 1-MS,RING leads roughly to the same performance of the 1-MS,ACT configuration. This means that the deployment of a tree-based passive architecture can lead to savings in terms of cost and energy with respect to a metro ring topology at the cost of worsening the performance in terms of blocking probability. On the other hand, simply adding a mesh links (1-MS + 1-MESH,ACT) to a tree-based active architecture allows to overcome the performance of the metro ring architecture.

We then consider the 2-MS configuration. Results are shown in the right-hand side of Fig. 9. Considering the tree-based active architecture, adding a mesh link between the two MSs (2-MS + 1-MESH, ACT) does not introduce

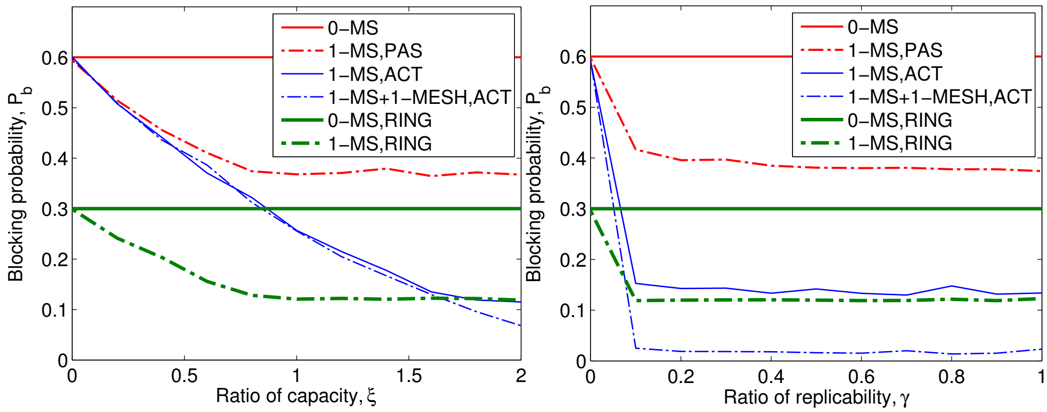

$0 \leq \xi < \infty$, $\gamma = 1$, $\rho = 2.5$ $\qquad$ $\xi = \infty$, $0 \leq \gamma \leq 1$, $\rho = 2.5$

Figure 10: Blocking probability for capacity and replicability constrained configurations in a realistic scenario

significant benefits with respect to the configuration without any mesh link (2-MS, ACT) since the amount of traffic that uses the mesh link is negligible. On the contrary, the addition of a second and especially of a third mesh link (2-MS + 2/3-MESH, ACT), can significantly decrease the blocking probability of the VoD requests. In fact, with respect to the 1-MS configuration, in this case the addition of these mesh links enables a more effective traffic balance between the two MSs. If we compare these results with the result obtained in case of the 2-MS,RING metro ring architecture, we can see how this configuration always overcomes the performance of the tree-based passive architecture. In addition, for values of $\rho > 2.5$, the performance is the same as for the 2-MS,ACT and 2-MS + 1-MESH,ACT configurations, but worse than the 2-MS + 2/3-MESH,ACT configurations. Instead, for values of $\rho < 2.5$ the 2-MS,RING configuration leads to better performance than the 2-MS,ACT and the 2-MS + 1/2-MESH,ACT configurations, but worse performance than the 2-MS + 3-MESH,ACT. This means that, in general, only the 2-MS + 3-MESH, ACT configuration, i.e., adding mesh links between all the nodes at the stage where the MSs are placed, guarantees better performance than the metro ring architecture in every traffic condition. In general, our simulations confirm that the VoD service performs better when an higher number of MS is deployed.

If we consider the capacity constraint for the ξ-bus of the MSs, results are shown in the left-hand side of Fig. 10. As ξ increases, a linear decrease of the blocking probability could be observed until a knee point is reached. After the knee point, the blocking probability does not decrease anymore. Different configurations yield to different values of ξ as knee point, as well as different minimum values of blocking probability. The value of the knee point, in each configuration, allows us to dimension the minimum capacity for the bus that leads to the minimum value for P_b. For example, in the 1-MS, PAS configuration, the knee point is around $\xi = 0.8$. This means that a bus capacity of 800 Mbit/s is sufficient to obtain the minimum possible value for P_b. Analogous considerations can be done while considering configurations with $m > 1$.

In the right-hand side of Fig. 10, we show the effects of the replicability γ ($0 \leq \gamma \leq 1$) when the traffic load is fixed at $\rho = 2.5$. Note that replicating less than 10% of the contents stored by the CDS-1 allows to reach the minimum blocking probability. This behavior is different than the behavior shown in Fig. 8 because here we consider content distributed according to a Zipf distribution ($\phi = 1.5$) and not to a uniform distribution. Under this assumption, the largest part of the requests is associated to the most popular contents. This is true for every configuration, also considering $m > 1$. Note also that the ϕ parameter of the Zipf distribution (see Section 4.1) is very important concerning content replicability, and thus a sensitivity analysis for this value is needed. A lower value of ϕ means that a smaller number of VoD requests is associated to the N most popular contents. The left-hand side of Fig. 11 shows the effect of a variation in the ϕ parameter on the blocking probability, considering the 1 MS,ACT configuration. For lower values of ϕ the blocking probability is higher, especially when we focus on lower values of γ. This happens because for small values of ϕ more VoD requests are associated to less popular contents that are not replicated in the MS and, thus, higher traffic will cross the bottleneck links.

6.3. Impact of future technologies

We finally evaluate the impact of future technologies on the results obtained so far. We assume that in the near future technological advances will provide to the users an access rate of 100 Mbit/s (instead of 10 Mbit/s), aggregation links of 1 Gb/s (instead of 100 Mb/s) and upper layer links of 10 Gbit/s (instead of 1 Gbit/s), to support video streams at much higher resolutions, such as 2K and 4K (see Table 4). The right-hand side of Fig. 11 shows

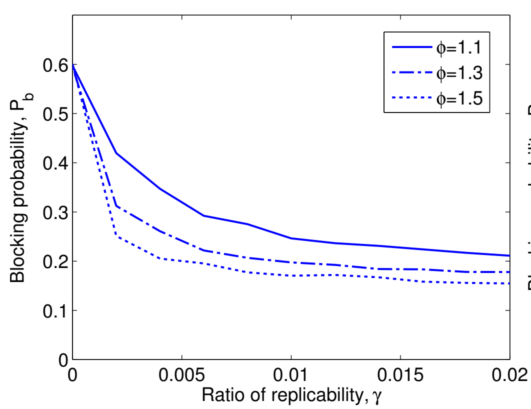

$\xi = \infty$, $0 \leq \gamma \leq 0.02$, $\rho = 2.5$, 1-MS,ACT

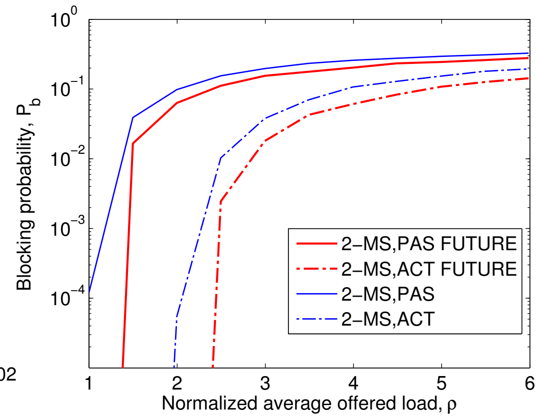

$\xi = \infty$, $\gamma = 1$, 2-MS

Figure 11: Sensitivity analysis for the ϕ parameter of the Zipf function for content replicability (*left figure*) and blocking probability comparison for the 2-MS configuration between the current scenario and the future-proof scenario (*right figure*)

Quality	Video	Audio	Bitrate
HD (720p)	1280x720 px	Dolby E-AC-3	7.5 Mbit/s
HD (1080p)	1920x1080 px	Dolby E-AC-3	12 Mbit/s
2K	2048x1080 px	Dolby E-AC-3	24 Mbit/s
4K	4096x2160 px	Dolby E-AC-3	60 Mbit/s

Table 4: Future bitrates for different video streaming settings

the comparison of blocking probability of such *future-proof scenario* with respect to the current scenario in terms of links capacity and video-streaming bitrates. We focus on the 2-MS unconstrained configuration ($\gamma = 1, \xi = \infty$) for both the passive and active tree-based architectures. Blocking probability in the future-proof scenario (2-MS,PAS FUTURE and 2-MS,ACT FUTURE) is slightly less than blocking probability in the current scenario (2-MS,PAS and 2-MS,ACT). These results point out how an effort by network operators aimed at increasing the users access bandwidth is needed in order to guarantee good performance for high-quality video streaming.

7. Conclusion

In this work we evaluated the performance of VoD delivery over different metro/access network architectures with replicated Metro Servers. First of all, we described a metro ring network architecture and two tree-based network architectures, one passive and one active, whose nodes can be equipped with MSs. We showed how the placement of MSs in different positions of such network architectures allow network operators to significantly decrease blocking probability of VoD requests. Moreover, the performance of the tree-based active architecture can be consistently improved by adding mesh links. The mesh link addition in the tree-based active architecture allows to gain better performance than the tree-based passive and especially than the metro ring architecture. On the contrary, the tree-based passive architecture is the most beneficial from a cost and energy consumption perspective, but it leads to the worst performance. We explored also the role of bus capacity and content replicability constraints for the MSs: in our condition of popularity, only a 10% of the contents should be replicated in the MSs in order to reach the minimum blocking probability for each network configuration. Moreover, savings can be achieved also in the bus capacity of the MSs: network operators, for each configuration, can dimension the minimum bus capacity to obtain the best performance. Finally, several issues still remain open for future research. For instance, dynamic scenarios for VoD traffic distribution can be investigated as well as different routing mechanisms, possibly taking into account energy aspects and the possibility to perform partial content replication, where only some excerpts of the content are replicated. Then, a detailed cost and energy analysis should be performed in order to comprehensively compare all the discussed architectures.

Acknowledgment

The research leading to these results has received funding from the European Community Seventh Framework Programme FP7/2013-2015 under grant agreement no. 317762 COMBO project.

[1] Cisco visual networking index: Forecast and methodology (2014-2019), 2015.

[2] R. F. Moghaddam, Y. Lemieux, M. Cheriet, 40 gbps access for metro networks: Implications in terms of sustainability and innovation from an LCA perspective, in: ICT4S'15 Conference, 2015.

[3] B. Qazzaz, R. Suppi, F. Cores, A. Ripoll, P. Hernandez, E. Luque, Providing interactive video on demand services in distributed architecture, in: Euromicro Conference, 2003.

[4] E. Di Pascale, D. Payne, M. Ruffini, Bandwidth and energy savings of locality-aware P2P content distribution in next-generation PONs, in: Conference on Optical Network Design and Modeling (ONDM), 2012.

[5] M. Mahloo, J. Chen, L. Wosinska, PON versus AON: Which is the best solution to offload core network by peer-to-peer traffic localization, in: Optical Switching and Networking, Vol. 15, 2015, pp. 1 – 9.

[6] R. Davey, D. Grossman, M. Rasztovits-Wiech, D. Payne, D. Nesset, A. Kelly, A. Rafel, S. Appathurai, S.-H. Yang, Long-reach passive optical networks, in: IEEE/OSA Journal of Lightwave Technology, Vol. 27, 2009, pp. 273–291.

[7] S. Lambert, B. Lannoo, A. Dixit, D. Colle, M. Pickavet, J. Montalvo, J. Torrijos, P. Vetter, Energy efficiency analysis of high speed triple-play services in next-generation PON deployments, in: Computer Networks, Vol. 78, 2015, pp. 68 – 82, special Issue: Green Communications.

[8] U. Mandal, C. Lange, A. Gladisch, P. Chowdhury, B. Mukherjee, Energy-efficient content distribution over telecom network infrastructure, in: International Conference on Transparent Optical Networks (ICTON), 2011.

[9] J. Baliga, R. Ayre, K. Hinton, R. S. Tucker, Architectures for energy-efficient IPTV networks, in: Optical Fiber Communication Conference and National Fiber Optic Engineers Conference, Optical Society of America, 2009.

[10] J. Araujo, F. Giroire, Y. Liu, R. Modrzejewski, J. Moulierac, Energy efficient content distribution, in: IEEE International Conference on Communications (ICC), 2013.

[11] E. Di Pascale, D. Payne, M. Ruffini, Impact of popularity evolution on P2P-based VoD delivery over next-generation optical access networks, in: IEEE Global Communications Conference (GLOBECOM), 2013.

[12] C. Jayasundara, A. Nirmalathas, E. Wong, C. A. Chan, Energy efficient content distribution for VoD services, in: Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011.

[13] G. Ciccarella, D. Roffinella, M. Vari, F. Vatalaro, Performance improvement and network TCO reduction by optimal deployment of caching, in: Euro Med Telco Conference (EMTC), 2014.

[14] A. Reaz, V. Ramamurthi, M. Tornatore, Cloud-over-WOBAN (CoW): An offloading-enabled access network design, in: IEEE International Conference on Communications (ICC), 2011.

[15] B. Tan, L. Massoulié, Optimal content placement for peer-to-peer video-on-demand systems, in: IEEE/ACM Transactions on Networking, Vol. 21, 2013, pp. 566–579.

[16] M. Broadbent, P. Georgopoulos, V. Kotronis, B. Plattner, N. Race, OpenCache: Leveraging SDN to demonstrate a customisable and configurable cache, in: IEEE Conference on Computer Communications Workshops (INFOCOM Workshops), 2014.

[17] P. Georgopoulos, M. Broadbent, B. Plattner, N. Race, Cache as a service: Leveraging SDN to efficiently and transparently support video-on-demand on the last mile, in: International Conference on Computer Communication and Networks (ICCCN), 2014.

[18] S. Kuenzer, J. Martins, M. Ahmed, F. Huici, Towards minimalistic, virtualized content caches with minicache, in: ACM Workshop on Hot Topics in Middleboxes and Network Function Virtualization, 2013.

[19] A. Shami, X. Bai, N. Ghani, C. Assi, H. Mouftah, QoS control schemes for two-stage ethernet passive optical access networks, in: IEEE Journal on Selected Areas in Communications (JSAC), Vol. 23, 2005, pp. 1467–1478.

[20] A. F. Milton, Star coupler for single mode fiber communication systems, 1976, US Patent 3,937,557.

[21] R. Fratini, M. Savi, G. Verticale, M. Tornatore, Using replicated video servers for VoD traffic offloading in integrated metro/access networks, in: IEEE International Conference on Communications (ICC), 2014.

[22] W. V. Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, P. Demeester, Trends in worldwide ICT electricity consumption from 2007 to 2012, in: Computer Communications, Vol. 50, 2014, pp. 64 – 76.

[23] U.S. Energy Information Administration, Average retail price of electricity to ultimate customers by end-use sector, 2015.
URL http://www.eia.gov/electricity/monthly/pdf/epm.pdf

[24] M. Savi, G. Verticale, M. Tornatore, A. Pattavina, Energy-efficient VoD content delivery and replication in integrated metro/access networks, in: IEEE Latin-American Conference on Communications (LATINCOM), 2014.

[25] W. Tang, Y. Fu, L. Cherkasova, A. Vahdat, Medisyn: a synthetic streaming media service workload generator, in: ACM International workshop on network and operating systems support for digital audio and video (NOSSADV), 2003.

[26] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, Q. Zhao, Modeling user activities in a large IPTV system, in: ACM SIGCOMM Conference on Internet measurement, 2009.

[27] P. Seeling, F. Fitzek, M. Reisslein, Video Traces for Network Performance Evaluation: A Comprehensive Overview and Guide on Video Traces and Their Utilization in Networking Research, Springer London, Limited, 2007.

[28] V. B. Iversen, et al., Teletraffic engineering handbook, Vol. 2, 2005.