Using Replicated Video Servers for VoD Traffic Offloading in Integrated Metro/Access Networks

Roberto Fratini, Marco Savi, Giacomo Verticale, Massimo Tornatore

Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, P.zza Leonardo da Vinci 32, Milano, Italy

roberto.fratini@mail.polimi.it, {savi,vertical,tornator}@elet.polimi.it

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/content delivery, network operators are currently considering novel integrated metro/access networks which accommodate replicated video servers directly in their infrastructure. In such way, servers for VoD delivery are placed nearer to the end users, the core segment of the network is partially traffic offloaded, and the end users experience better performance in terms of QoS. In our work, we will evaluate the performance improvement of an integrated metro/access architecture for VoD delivery with replicated video servers considering different configurations in terms of number of replicated servers, meshing degree and adopted network technologies. We develop a network simulator in which replicas of video servers (called Metro Servers, or MSs) are deployed to meet the demand of VoD traffic. In the result section we compare the performance of the various configurations and discuss which are the minimum requirements to minimize blocking of the VoD requests.

I. Introduction

Due to the rapid adoption of new broadband services such as Media Streaming, File Sharing, VoIP, and Online Gaming, Internet traffic keeps increasing steadily. According to a recent forecast report by Cisco [1], Internet media streaming traffic will soon become the dominant type of Internet traffic, and especially Video-on-Demand (VoD) has been recognized as one of the most prominent media streaming services. To sustain the growth of these broadband services, new optical access network technologies, such as, e.g., fiber-to-the-home (FTTH) or fiber-to-the-building (FTTB), are being deployed in order to provide fiber connectivity closer to the end users and higher user access rates. Such increase of user access rates will pose a serious scalability issue to the current telecommunication network infrastructure.

Current telecommunication networks consist of three domains: access, metro/aggregation, core. Core network provides nationwide or global coverage, and spans long distances, i.e., each link (using optical fibers) is 100s to 1000s of km in length. Metro network typically spans a metropolitan region. Access network connects end users to their service provider, enables end-users to connect to the rest of the network infrastructure, and spans a distance of a few km.

Currently, most video contents are stored in centralized media servers placed in the core segment of the network and, without new solutions for content delivery, the metro and core segments of the network will soon be flooded by media traffic: various studies have identified the risk of them constituting a possible traffic bottleneck that could severely compromise the performance of VoD services.

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

Nowadays 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) [2], represent a promising candidate for next-generation cost- and energy-efficient network architectures for content delivery to users. They are typically multi-stage networks spanning up to few hundreds of kilometers and, because of their extended footprint, the effect of replicated servers in this context needs to be investigated. While some studies have already investigated the effect of positioning server replicas (caches) in core nodes [3] and/or in access nodes [4], [5], [6], to the best of our knowledge, the effect of VoD content migration/distribution at the various stages of this novel integrated metro/access network hierarchy has never been investigated.

In this work, we will study the performance of VoD content delivery over an integrated metro/access network architecture, where Metro Servers (MSs) are placed at different stages of the metro/access network. We will show how different metro/access network topologies, equipped with different technologies and under different configurations, perform in terms of blocking probability for VoD requests.

The rest of the paper is organized as follows: in Section II we provide an overview of related works. In Section III we describe some possible representative architectures for an integrated metro/access network architecture and we show at which hierarchical level the replicated Metro Servers can be placed to enable offloading of VoD streaming traffic. In Section IV we focus on the VoD traffic modelling and content generation, explaining our assumptions and how the system has been simulated. In Section V we quickly overview the
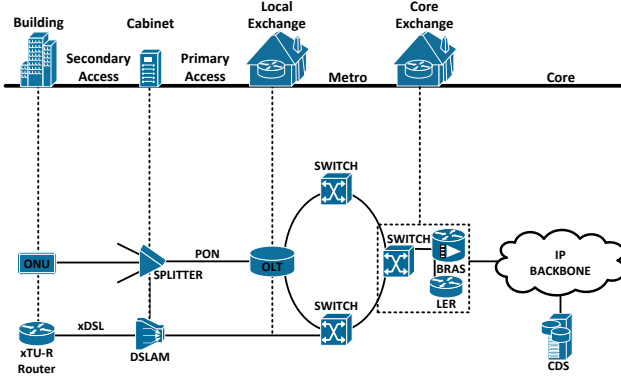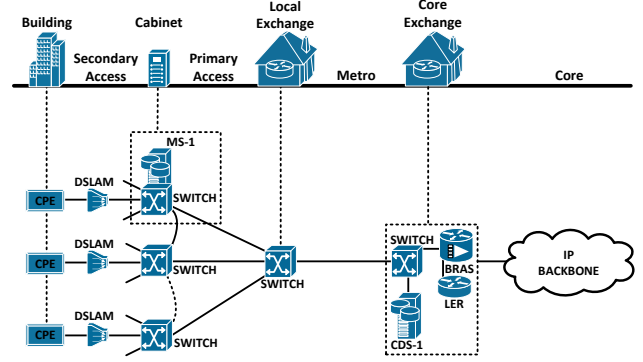
Figure 1: Today's network architecture

Figure 2: Proposed electrical integrated metro/access network architecture

Figure 3: Proposed optical integrated metro/access network architecture

main features of our simulator. Then, Section VI shows and discusses the main results of our simulations. Finally, Section VII draws the conclusion of this study.

II. RELATED WORKS

Some existing studies have dealt with performance evaluation and optimization of content placement and distribution in telecommunication networks. One of the main challenges consists in placing the content servers taking into account jointly energy efficiency, core-traffic offloading and latency. Ref. [7] analyzes the performance of a centralized VoD distribution system where a single replicated server is placed in different segments of the network (access/metro/core), concluding that placing the replicated server in the metro segment is the most energy-efficient solution if it is possible to switch it off during off-peak hours of service. Ref. [4], Ref. [5] and Ref. [6] investigate the performance of a distributed system of caching servers located specifically in the access network. Ref. [4] and Ref. [5] propose to adopt peer-to-peer (P2P) solutions, where contents are stored and shared by end users in a LR-PON underlying architecture. Ref. [6] defines an hybrid wireless/optical access network and studies a cost-efficient solution in cashing servers placement over such network. Performance of a distributed system of caching servers located in the core network is analyzed in Ref. [3].

In our work we will study the performance of VoD content placement at various levels of a novel integrated metro/access network hierarchy, which is a dimension of possible replicated server placement that has never been explored before.

III. METRO/ACCESS NETWORK ARCHITECTURES FOR VoD DELIVERY WITH REPLICATED SERVERS

In this Section, we first describe the generic architecture of current metro/access networks. Then we introduce the concept of *integrated metro/access network architecture* and we differentiate between one architecture based mostly on electrical (e.g., Ethernet) switches and one architecture based mostly on passive optical equipments. For both solutions, we discuss how replicated servers can be introduced in the architecture (metro/core offloading architecture).
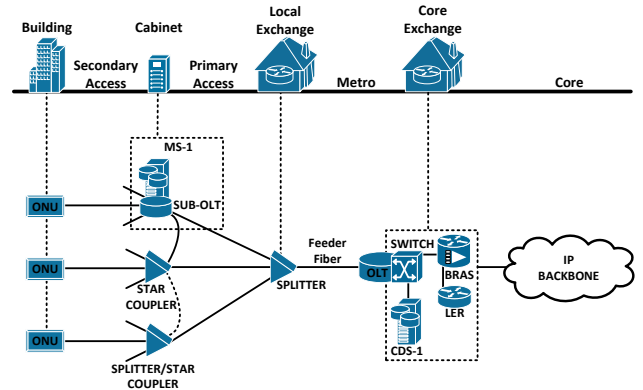
A. Today's network architecture

As shown in Fig. 1, today's metro network architectures are mainly characterized by a metro ring (e.g., using Gigabit Ethernet) which interconnects different access technologies, such as Digital Subscriber Line (xDSL) 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, which typically hosts the content delivery servers (CDSs).

B. Integrated metro/access architecture

We consider here a new integrated metro/access architecture where MSs can be placed at various stages of the metro/access network's nodes. In this architecture the metro/access network is characterized by a long-reach branch-tree architecture.

1) Metro Servers and additional mesh links placement: In such integrated scenario, 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) in order 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 of
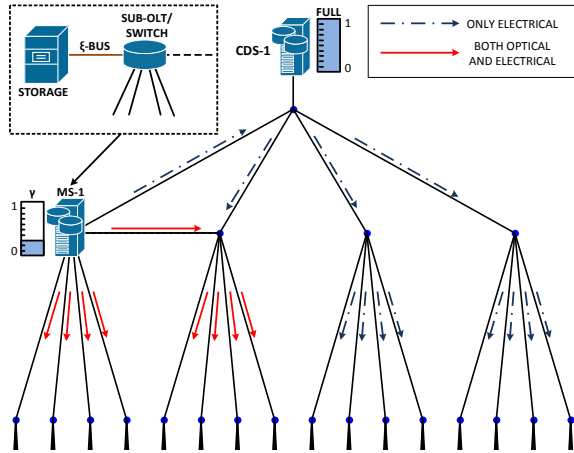
Figure 4: Example of a 64-users hybrid meshed/branch-tree metro/access network (1-MS + 1-MESH configuration).

the architecture allow to increase the number of end users that can take advantage of that specific MS.

In our model, we assume that a primary content server (referred to as CDS-1, and storing all the contents) is placed somewhere in the core network (e.g, at the core exchange, as in Fig. 2 and Fig. 3).[1] This case with only CDS-1 is referred as a 0-MS configuration.

2) Electrical vs. optical architecture: As shown in Fig. 2 and Fig. 3, the integrated network described above might be implemented using either optical or electrical technology. In the electrical case, the nodes of the hybrid meshed/branch tree are Ethernet switches and the addition of DSLAMs and Customer Premises Equipments (CPEs) provide broadband access to the end users. In the optical case, an Optical Line Termination (OLT) in the core exchange performs electrical/optical conversion for the downstream traffic, the intermediate metro nodes of the hybrid meshed/branch tree 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 of the two solutions, and the presence of replicated servers plays a relevant role in choosing the most effective option. *i*) To begin with, the electrical solution can more easily incorporate the introduction of an MS and/or of a mesh link in the intermediate stages of this metro/access network. In fact, the electrical architecture does not require specific upgrades to support the introduction of a MS or of a mesh link, while in the optical architecture some significant changes in terms of network architecture and control are required. Specifically, in the optical case, whenever an MS is introduced in an intermediate metro node, the node is no more passive and an additional active component is needed, i.e. a sub-OLT

(see [8] for more details). Such active device works as an 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 traffic which is directed to the sub-OLT's child nodes, allowing content delivery from the MS to the child nodes of the node where the MS (and, consequently, the sub-OLT) is placed. Moreover, the addition of a mesh link between the MS and an adjacent node requires the substitution, in the adjacent node, of the power splitter with a star coupler [9], which is a passive device with more than one input (see Fig. 3).[2] *ii*) On the other hand, if we consider cost and energy efficiency, the optical architecture might represent a more desirable solution than the electrical one. In fact, in the optical 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 electrical 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. *iii*) Finally, we must also consider that an electrical solution allows more users to be served by a single MS. In fact, switches in an electrical 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. E.g., we can see in the example reported in Fig. 4 (where red and dotted blue arrows depict the possible directions of flows from the MS to end users, in the optical and in the electrical case, respectively) that the amount of users reachable in the electrical case is much larger than in the optical case.

In the next Section, we will define a practical modelling of the VoD traffic served over these network architectures, which is our reference model for the following performance evaluation.

IV. VoD CONTENT AND TRAFFIC MODELLING

Different models for VoD contents and traffic can be used. In this paper we use the results in [10], where a long-term trace of a real VoD streaming traffic is collected and statistically analyzed to model the following properties of a VoD system.

A. Connection request

According to [10], the probability distribution function of video-requests can be modelled by a Poisson distribution. For every hour of the day, this distribution is characterized by a different λ according to a diurnal access pattern.

---

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

[2]In this case, an intermediate node might receive traffic both from the root OLT and from the sub-OLT, so, to avoid collision of traffic from these different sources, a coordination protocol between the OLT and the sub-OLT is needed. A generalization of this scheme is possible also for multiple MSs and mesh links. Discussion about the OLT-subOLTs coordination protocol and about a detailed implementation of such an architecture is out of the scope of this study.
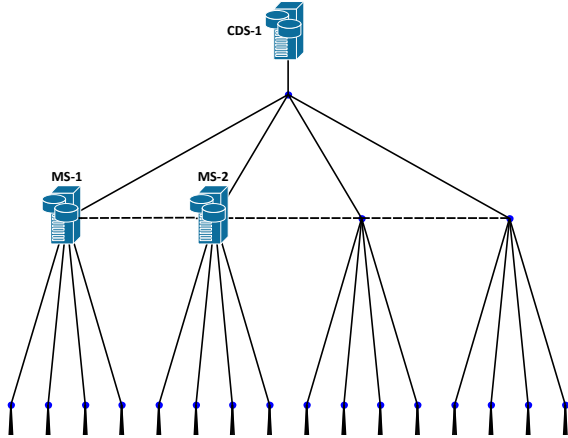
Figure 5: Example of a 2-MS + n-MESH configuration, with $n = 3$

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 I: Bitrates for different video streaming settings

B. Content duration

Content duration distribution is largely determined by the nature of the contents which are stored in the content servers, i.e. short/long clips [10], [11]. Thus, content duration distribution results in a multi-modal distribution having multiple peaks representing different trends of durations. In our work a finite mixture-gaussian as defined in [10] is used in order to capture different peaks of the duration distribution for short/long clips.

C. Content popularity

Video popularity is statistically modelled through the Zipf distribution [12], which is often used for content popularity of large VoD media systems (e.g., VoD streaming of an ISP).

D. Encoding bitrate

A set of common used streaming bitrates is defined in order to model a discrete distribution of constant bitrates, as reported in Table I.

E. 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 a MS (see Fig. 4). The most popular contents, according to the Zipf distribution, are replicated first.

F. Capacity constraint

Each MS is connected to the TX device (i.e., a switch in the electrical architecture and a sub-OLT in the optical architecture) 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 the network links, that is supposed to be the same for all the links:

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

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.

V. 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 circuit between a content server and an user has to be established for each VoD request. Whenever a VoD request is generated by an user, it is met by an MS or by the CDS-1 following the *nearest routing algorithm*. The nearest routing algorithm checks the MSs that store the requested content starting from the nearest ones 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, bandwidth for each link of the path is allocated and the MS is marked as source for the current request (i.e., the MS is available). Otherwise, another path between the selected MS and the user, if any, is considered. If no path can allocate resources for the current request, the selected MS is considered unavailable and the following one is investigated. If no MS is available, the request is forwarded to the CDS-1. If the CDS-1 is unavailable, the request is discarded (i.e., blocked).

For our simulations, we adopt the content and traffic modelling we defined in Section IV.[3]

VI. PERFORMANCE EVALUATION

We consider the scenario depicted both in Fig. 4 and in Fig. 5. We simulate a generic hybrid meshed/branch tree metro/access network, with a per stage splitting ratio of 1:4. The overall network consists of three stages leading to a total splitting ratio of 1:64 (64-users network model). The capacity for each link is $C_{link} = 1$ Gbit/s. For this specific topology, we evaluate different configurations in terms of the adopted technology (i.e., optical or electrical), the number of MSs (1-MS and 2-MS) and the mesh links added (1-MESH, 2-MESH and 3-MESH).

A. 1-MS configuration

The 1-MS configuration is characterized by a single MS placed in the second splitting stage of the long-reach integrated metro/access network (Fig. 4). The configuration can be improved increasingly adding up to 3 mesh links (1,2,3-MESH configurations). Blocking probability P_b related to different network configurations is shown in Fig. 6. In this scenario, we assume that the bus capacity and the storage dimension of the

---

[3]A theoretical validation of our simulator, under simplified assumptions, has been performed and results of the simulations match results obtained from a theoretical analysis. Details of the theoretical validation are not reported for sake of conciseness.
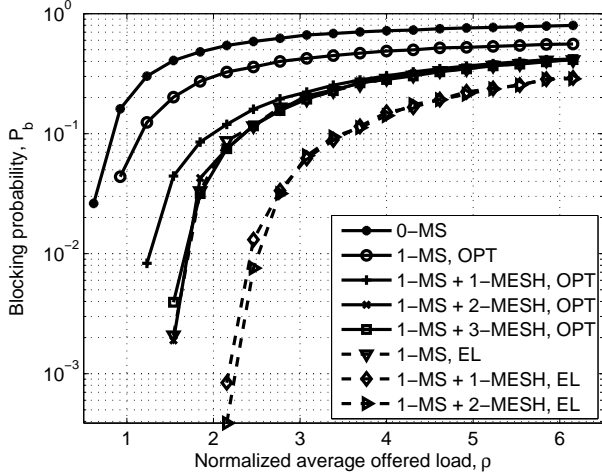
Figure 6: Blocking probability for increasing load for the 1-MS configuration in a 64-users network ($\xi = \infty$, $\gamma = 1$)
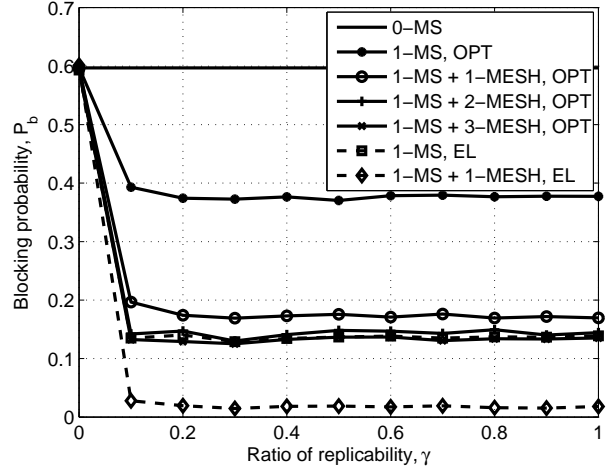

Figure 8: Blocking probability for increasing ratio of replicability γ in a 64-users network ($\rho = 2.5$, $\xi = \infty$)
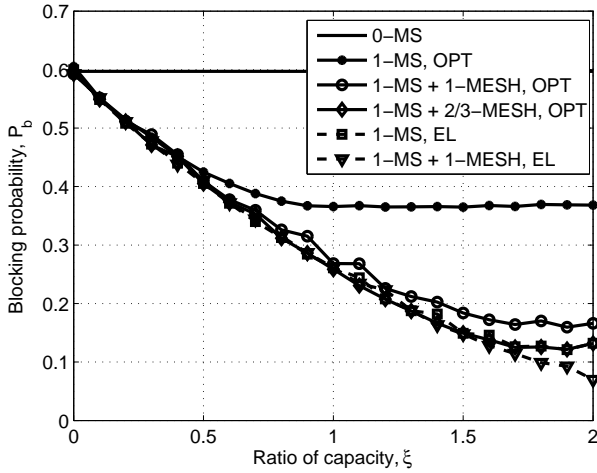

Figure 7: Blocking probability for increasing ratio of capacity ξ in a 64-users network ($\rho = 2.5$, $\gamma = 1$)

MS are large enough to prevent it to be a bottleneck element of the network architecture ($\xi = \infty$, $\gamma = 1$). We can see how in 1-MS, OPT and 1-MS + 1-MESH, OPT configurations, P_b decreases more than an order of magnitude with respect to the 0-MS case, while adding more than one meshing links does not bring significant benefits, especially at high traffic loads. In addition to this, we might observe how P_b significantly decreases when we consider the electrical architecture respect to the same configuration in the optical architecture. In fact, switching devices in a 1-MS electrical architecture (1-MS, EL), as discussed in Section III, respect to the 1-MS, OPT case enable traffic also in the upstream direction and even just a single MS can reach all the users. The same effect of the 1-MS, EL case, as we can see, can be achieved in the optical architecture by placing 3 mesh links over the optical network: the 1-MS + 3-MESH, OPT optical case perfectly

matches the 1-MS, EL electrical case. A further significant decrease of P_b is achieved in the 1-MS + 1-MESH, EL 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, only marginal improvements are achieved at increasing the number of mesh links in the electrical case.

If we consider the capacity constraint for the ξ-bus of the MS, results are shown in Fig. 7. We plot the blocking probability as a function of the ratio of capacity ξ, with the traffic load fixed at $\rho = 2.5$ and unconstrained ratio of replicability ($\gamma = 1$). Note that when $\xi = 0$ all the configurations return the same blocking probability since $\xi = 0$ is equivalent to the 0-MS configuration. As ξ increases, a linear decrease of 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, OPT 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.

In Fig. 8, we measure the effects of the replicability γ ($0 \leq \gamma \leq 1$) in the configurations discussed so far, with the traffic load fixed at $\rho = 2.5$ and unconstrained ratio of capacity ($\xi = \infty$). Note that replicating only a 10% of the contents stored by the CDS-1 allows to reach the minimum blocking probability. This behaviour comes from the Zipf distribution, as the largest part of the requests ($\sim$90%) is for the most popular contents.

B. 2-MS configuration

The 2-MS configuration is characterized by the placement of two metro-servers in the metro/access section of the network. Several placements configurations are possible. We choose to
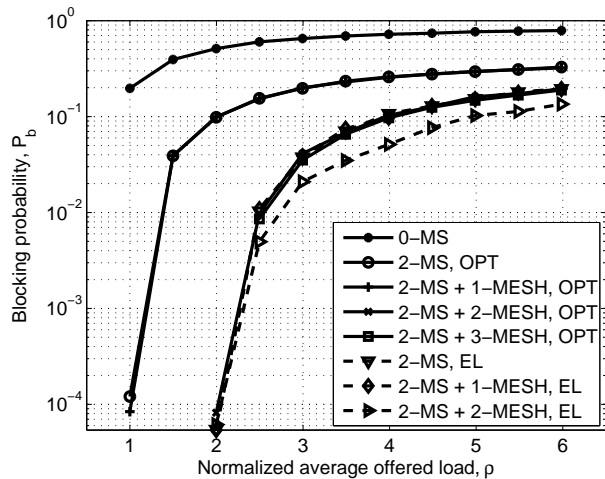
Figure 9: Blocking probability for increasing load for the 2-MS configuration in a 64-users network ($\xi = \infty$, $\gamma = 1$)

place the MSs in the first and in the second node of the second splitting stage, as shown in Fig. 5.

Results are shown in Fig. 9. We assume that the bus capacity and the storage dimension of the MSs are large enough to prevent it to be a bottleneck element of the network architecture ($\xi = \infty$, $\gamma = 1$). We can see how the 2-MS, OPT line in Fig. 9 is closely similar to the 1-MS + 1-MESH, OPT curve of Fig. 6. Moreover, placing an additional meshing link on the 2-MS configuration (2-MS + 1-MESH, OPT) between the two MSs does not introduce significant benefits since the amount of traffic that uses such link is very low. Then, adding a second and a third mesh link (2-MS + 2-MESH, OPT and 2-MS + 3-MESH, OPT), decreases the blocking probability of the network as expected. In particular, the performance of these configurations is strictly comparable to the performance of the 2-MS, EL and 2-MS + 1-MESH, EL. In addition, the configuration 2-MS + 2-MESH, EL further decreases blocking probability, reaching the best possible performance.

VII. Conclusion

In this work we evaluated performance of VoD delivery over a metro/access network with replicated servers (i.e., Metro Servers, MSs). First of all, we defined two integrated metro/access network architectures, one optical and one electrical, whose nodes can be equipped with MSs. Then, we designed and developed an event-driven discrete-event simulator by which we modelled VoD delivery over the two network architectures. We have shown how the placement of MSs and mesh links in strategical positions of the integrated metro/access network allows network operators to significantly decrease blocking probability of VoD requests in both optical and electrical network architectures. In particular, considering the same topology, the electrical architecture allows to obtain better performance than the optical one, but it is less cost- and energy-efficient.

We explored also the role of bus capacity and replicability content constraints: 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, by introducing the possibility to dynamically create contents and allow them to have dynamic popularity. Moreover, different routing mechanisms should be explored in order to discover new algorithms which optimize network performance.

References

[1] "Cisco visual networking index: Global mobile data traffic forecast update, 2012-2017," February 2013.

[2] R. P. Davey, B. Grossman, M. Rasztovits-Wiech, D. B. Payne, D. Nesset, A. Kelly, A. Rafel, S. Appathurai, and S.-H. Yang, "Long-reach passive optical networks," *IEEE Journal of Lightwave Technology*, vol. 27, no. 1, p. 273, 2009.

[3] U. Mandal, C. Lange, A. Gladisch, P. Chowdhury, and B. Mukherjee, "Energy-efficient content distribution over telecom network infrastructure," in *Transparent Optical Networks (ICTON), 2011 13th International Conference on*, 2011, pp. 1–4.

[4] E. Di Pascale, D. Payne, and M. Ruffini, "Bandwidth and energy savings of locality-aware P2P content distribution in next-generation PONs," in *Optical Network Design and Modeling (ONDM), 2012 16th International Conference on*, 2012, pp. 1–6.

[5] M. Ruffini, E. Di Pascale, and D. Payne, "Improving high fidelity multimedia distribution in next-generation optical networks," in *Proceedings of ICTON 2013*, 2013.

[6] A. Reaz, V. Ramamurthi, and M. Tornatore, "Cloud-over-WOBAN (CoW): An offloading-enabled access network design," in *Communications (ICC), 2011 IEEE International Conference on*, 2011, pp. 1–5.

[7] C. Jayasundara, A. Nirmalathas, E. Wong, and C. A. Chan, "Energy efficient content distribution for VoD services," in *Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011 and the National Fiber Optic Engineers Conference*, 2011, pp. 1–3.

[8] A. Shami, X. Bai, N. Ghani, C. Assi, and H. Mouftah, "QoS control schemes for two-stage ethernet passive optical access networks," *Selected Areas in Communications, IEEE Journal on*, vol. 23, no. 8, pp. 1467–1478, 2005.

[9] A. F. Milton, "Star coupler for single mode fiber communication systems," Feb. 10 1976, uS Patent 3,937,557.

[10] W. Tang, Y. Fu, L. Cherkasova, and A. Vahdat, "Medisyn: a synthetic streaming media service workload generator," in *Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, ser. NOSSDAV '03. New York, NY, USA: ACM, 2003, pp. 12–21.

[11] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, and Q. Zhao, "Modeling user activities in a large IPTV system," in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, ser. IMC '09. New York, NY, USA: ACM, 2009, pp. 430–441.

[12] P. Seeling, F. Fitzek, and 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.