



UWL REPOSITORY

repository.uwl.ac.uk

Quality of service in an information-centric network

Al-Naday, Mays, Bontozoglou, Andreas, Vassilakis, Vassilios and Reed, Martin (2015) Quality of service in an information-centric network. In: IEEE Global Communications Conference, 08-12 Dec 2014, Austin, USA.

<http://dx.doi.org/10.1109/GLOCOM.2014.7037079>

This is the Accepted Version of the final output.

UWL repository link: <https://repository.uwl.ac.uk/id/eprint/2819/>

Alternative formats: If you require this document in an alternative format, please contact: open.research@uwl.ac.uk

Copyright:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy: If you believe that this document breaches copyright, please contact us at open.research@uwl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Quality of Service in an Information-Centric Network

Mays F. AL-Naday, Andreas Bontozoglou, Vassilios G. Vassilakis, Martin J. Reed

School of Computer Science and Electronic Engineering

University of Essex

Colchester, Essex CO4 3SQ, UK

Email: mfhaln@essex.ac.uk, bodozoglou@essex.ac.uk, vasilak@essex.ac.uk, mjreed@essex.ac.uk

Abstract—QoS provisioning is one of the key challenges facing current as well as future Internet architectures. Its dependency on content recognition does not allow a straightforward support of QoS in the IP, host-centric, model. In contrast, Information-Centric Networking (ICN) offers native content identification in the network, which can be exploited to develop a common, elegant, framework for supporting QoS-based delivery. Therefore, ICN may naturally overcome many of the cumbersome fixes and limitations of today's solutions. In this work, we exploit the flexibility in semantic representation offered by ICN to present a flexible and scalable ICN-based QoS model. Our model defines QoS requirements as information items that can be linked to the content at various aggregation levels, independent of the communication approach. Therefore, it can be applied uniformly to various network types and hierarchies. Furthermore, our model offers enhanced traffic treatment as well as resource utilization while significantly reducing the overhead on the network.

Index Terms—Information-Centric Networks, QoS, DiffServ, IntServ, MPLS-TE.

I. INTRODUCTION

Recent years have witnessed an active research in information-centric networking (ICN) architectures. However, QoS is one of the key challenges that has not been considered widely. Yet, content-based traffic have witnessed a rapid growth in terms of volume and heterogeneity as well as a broadening range of granularity. Such growth stresses the need for flexible, scalable and affordable QoS models that can be supported at different aggregation levels. Current QoS models suffer from severe drawbacks; for example, DiffServ is only suitable for applying at access networks, to map communication sessions into priority-based flows, while IntServ has significant scalability problems. In general, this is not adequate to identify the proper treatment of the carried content and, further, deep packet inspection is required to enforce traffic policing and shaping. This, in turn, raises concerns regarding the privacy and integrity of the carried content [1]. DiffServ also requires a state of “flows-to-class” mapping in each edge router, which may require high overhead. Although, this is limited to edge routers, to help with scalability in the core, this comes with the disadvantage that it is not feasible to then account for traffic further in the core based upon simple ingress policing. Consequently, the advantages of DiffServ over flat over-provisioning are not always clear [2]–[4]. IntServ secures the required resources for pre-defined traffic flows. However,

the cost of signaling and resource reservation overhead and per-flow resource maintenance significantly limit the scalability of IntServ, thereby constraining its usability to critical, delay sensitive, traffic flows of limited size and dissemination scale.

MPLS traffic engineering (MPLS-TE) introduces modified models of QoS, DiffServ TE (DS-TE), that can be offered in core networks to aggregated traffic flows [5]. This is achieved by allowing DiffServ classification to be applied to LSPs of pre-defined SLAs (i.e. equivalent to IP-based IntServ paths) [2], [6]. DS-TE based models offer better scalability than that of IP based DiffServ and IntServ; but they are still constrained by the overhead limitations of RSVP-TE. Therefore, in practice, their applicability is limited to highly aggregated flows, irrespective of the aggregation type. Furthermore, hierarchical networks require a cumbersome mapping between different DiffServ marking schemes, which introduces further complexity and processing overhead at the edge nodes.

The complexity, and hence the scalability limitations of IP and MPLS based QoS models is mainly derived by the lack of information awareness in the network. Many of these limitations can, naturally, be resolved in an ICN environment. This is due to the shift of the communication paradigm from the end-host to the information, where the network identifies the requested content before locating it. This shift allows the network to natively identify the advertised information, using a different addressing scheme from that of the network connectivity, and construct the communication relations based on common information interests. It further reduces, or eliminates, the need to perform deep packet inspection to identify content as it can be recognized by its own identifiers. Moreover, ICN, utilizing suitable delivery mechanisms, allows QoS models to be reflected directly on the network paths and offers natural QoS description, with negligible overhead, by utilising the information namespace.

Generally, ICN architectures [7]–[11] offer native content identification; thereby, considerably simplifying the process of mapping information to QoS requirements. However, they may support different QoS models with varying features, depending on the functional model of each architecture. For example, the PURSUIT architecture offers routing qualities that simplifies the provision of IntServ-like routes but with no signaling overhead, thereby introducing significant resource

saving and better scalability [12]. In this work, we exploit the advantages of source routing offered in ICN architectures, such as PURSUIT, to present a novel QoS model that facilitates a DiffServ type of integration over IntServ-like routes. In our model, delivery trees are constructed using virtual links that reflect the desired QoS requirements with no signaling overhead along the path and minimal state in the forwarding node.

The rest of the paper presents, in Section II, a discussion of QoS support in various ICN architectures with focus on semantic representation and routing mechanisms as key variables affecting QoS delivery. Section III presents our model and outlines an example set of the QoS delegations that can be supported for various traffic granularities as well as their facilitation using source routing mechanisms. Section IV presents an example implementation of our model in the PURSUIT testbed; while Section V delivers our evaluations, and finally Section VI draws our conclusions.

II. ICN ARCHITECTURES AND QoS

Generally, ICN architectures implicitly simplify the task of mapping content to QoS classes. This is achieved through the information identification functionality, thereby diminishing many of the privacy and integrity concerns raised by the current, cumbersome, traffic policing methods. Information recognition can either be offered by a native information management function as part of a clean-slate ICN architecture [7], [8], [11]; or via a content mediation layer as part of an overlay platform [10]. A comprehensive comparison of the developed ICN architectures is provided in [13], [14].

NDN [7] offers a hierarchical content naming that allows content-based aggregation mechanisms for various granularities, depending on the matching prefix, for example a content provider may dedicate a QoS-based prefix for all VoD content. However, as content grouping is solely provided by the publisher, it is difficult for third-parties, such as the network operator, to introduce alternative, quality-defining, aggregation groups. Therefore, a mapping and “type-defining” functions are required to determine the proper treatment of content. In addition, NDN delivery paths are provisioned on a hop-by-hop basis. Therefore, the architecture inherits the same resource management issues witnessed in IP networks.

COMET [10], [15] offers a block-based content identification structure where inter-related content items are identified with a single block of sequential identifiers to achieve an aggregate content resolution. Each identifier points to a registered *Content Record*, which maintains state of the known versions of a content item and the corresponding Class of Service (CoS) as well as the QoS constraints. This information is utilized for source selection and path provisioning. Signaling overhead is required to configure the nodes along a path; and a state of “content-to-path” mapping is to be maintained by the ingress/egress routers of each domain. Although, such state is kept at the edges of the domain, considering the growing size of the information space, it potentially suffers from scalability limitations.

PURSUIT [8], [9] offers an information management plane where information items are grouped within *scopes* that reflect common semantics. Scopes are classification-based information, linked in hierarchical structures to achieve information aggregation. Information items can, also, be associated with other items to reflect a desired relation. This plane allows QoS requirements to be expressed as information items, which can be advertised by any, authorized, party. Moreover, by utilizing the scoping hierarchy, QoS information can be aggregated within particular scopes and linked to information items at various aggregation levels. This is illustrated by the example presented in Section III-A. Furthermore, the source routing mechanism of the PURSUIT architecture [12], which encodes all the links of the path in the packet’s header, offers the advantage of resource reservation without the signaling overhead and scalability limitations of such mechanism as IP IntServ.

We utilize this ability to provide a class-based differentiation over the IntServ-like paths by introducing link virtualization to the network, whereby each physical link is divided into a number of virtual links. Each virtual link is identified with a separate, Bloom Filter (BF) based, *Link Identifier* (Lid) and corresponds to a predetermined CoS of specific constraints. We show that our approach can achieve a significant resource saving compared to DiffServ models in IP, and therefore significantly enhances the traffic admission rates due to the improved resource utilization.

III. QoS MODEL FOR THE PURSUIT ARCHITECTURE

The PURSUIT architecture defines the communication paradigm as one of a publish/subscribe nature, where information owners publish the availability of their information and interested parties subscribe to it. The term *information* does not only refer to content but rather covers a broad range of publications, such as: content, request for an action, notifications, etc. Information items are identified by flat, possibly unique, Rendezvous Identifiers (Rid) and can be affiliated to, potentially, any number of scopes, identified by Scope Identifiers (Sid). When a mutual interest in an item occurs, a pub/sub relation can be established. The architecture defines a functional model, consisting of three core functions, to enable the establishment of pub/sub relations: *Rendezvous* (RV): responsible for matching supply with demand, *Topology Management* (TM): calculates and constructs the delivery trees, represented by LIPSIN-like *Forwarding Identifiers* (Fid) [12] and *Forwarding* (FW): responsible for decoding the BF-based Fid using simple and fast bitwise operations to perform packet switching. The network connectivity is expressed by flat, BF-based LIDs where each edge in the network has at least two LIDs, one in each direction. Link virtualization can be achieved by allocating more than one Lid per edge direction as highlighted in the previous section.

A. Scoping Hierarchy To Support QoS

QoS in our functional model is concerned with how to satisfy different traffic requirements on a shared pool of resources using an ICN pub/sub communication approach. Two

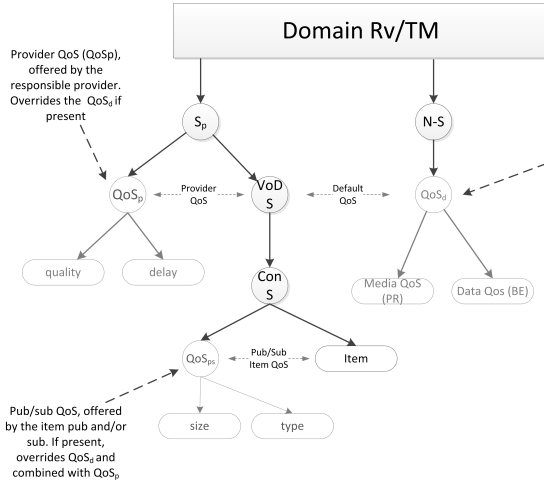


Fig. 1. A Sample QoS Scoping Hierarchy in the PURSUIT Information Management Plane Showing Three Potential Models of QoS Provisioning

examples of QoS delegation scenarios will be considered to show that the functional model provides a flexible framework:

- **Highly aggregated QoS (HA-QoS):** applies a single QoS strategy across a broad range of pub/sub relations. Traffic rates of such relations are aggregated and this aggregated rate is reserved in the available capacity in each forwarding node. It is up to the external network feeding the network ingress to check that each micro-relation can be supported in this aggregate QoS strategy. This is broadly comparable to the goals of MPLS-TE in contemporary IP/MPLS carrier grade networks.
- **Prioritized QoS (P-QoS):** allocates different priorities to different information items, a priority value reflects a set of delivery constraints. The traffic rate may not be allocated for each micro-relation but high priority traffic generally gets better treatment in the forwarding nodes. This is broadly comparable to the goals of IP DiffServ.

Although only two models of QoS are investigated in this paper, the architecture allows a very large range of QoS mechanisms to be implemented; for example, a fine-grained QoS could be implemented on each micro-relation, which resembles IP IntServ.

Figure 1 illustrates the QoS scope hierarchy for an example of the QoS scenarios. The domain RV offers a default HA-QoS scope, QoS_d , based on types of information items; for example, video items have different QoS metrics than data items. QoS_d could be used when no party (provider, publisher, subscriber) express specific QoS requirements. Alternatively, the responsible, content or network, provider may choose to enforce its own QoS strategy by creating an associated scope, QoS_p , to the content scope where it publishes the QoS information. The depth of QoS_p hierarchy is decided by the provider and could go as far as the item level; however, we assume that QoS_p is, generally, in the form of a HA-QoS. Offering QoS_p would effectively override QoS_d . Finer QoS, QoS_{ps} , can also be enforced from the publisher and/or subscriber. However, if QoS_p is already enforced by the

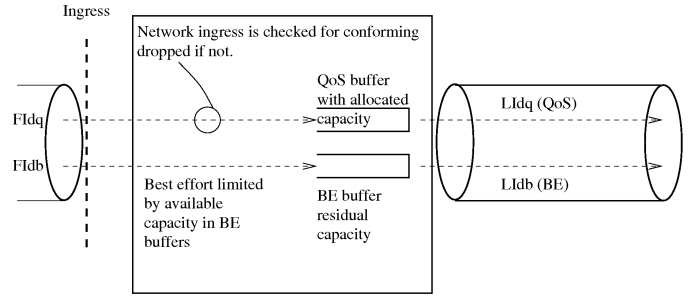


Fig. 2. QoS differentiation in the PURSUIT architecture showing an example of two virtual links: QoS-based with allocated capacity and Best Effort (BE) with residual capacity

provider then QoS_{ps} , in general, can still be offered as long as it does not violate QoS_p , otherwise the latter would override it. If both the publisher and subscriber publish their QoS requirements, a mediation mechanism would be applied to determine the final QoS. It should be noted that the content subscriber can, also, be a publisher of QoS information.

B. QoS and Source Routing

QoS differentiation can be provided by allocating multiple LIDs to each edge direction, where each LID corresponds to a certain buffer. In Fig. 2 an example of two LIDs per edge direction is shown: one corresponds to QoS traffic and one reflects Best Effort (BE) traffic. More buffers can be established to either support different priorities of traffic or different types of QoS. In simplistic terms, if the traffic rate of all of the QoS traffic passing through a forwarder is known then it can be allocated to a QoS buffer that has a non-blocking forwarding rate to the output link equal to, or more than, the total QoS traffic rate. In practice this is more complex as there is no concept of an exact instantaneous traffic rate in packet networks; instead the rate is generally described as a token bucket (or leaky bucket) using the tuple $\langle \text{token rate, token depth} \rangle$. In addition to the rate there may be requirements on minimum forwarding delay or minimum delay variance (jitter). This range of information is generally enough to allocate a pub/sub relation to a buffer.

Consider first the P-QoS strategy. In IP, and generally IP/MPLS, this requires a field in the header which in pure IP is the type of service and precedence bits. However, in the ICN version as shown in Fig. 2 there is no need for the addition of option fields, the priority can be inferred directly from the FID. Consider now the HA-QoS or FG-QoS strategies. These require traffic to be allocated to particular buffers, and optionally may require that the buffer parameters are updated to allow for the extra traffic. The functional model and LIPSIN forwarding allow a unique pub/sub relation to be allocated to a specific QoS path (and associated buffers) using a specific FID. This, allows QoS to be implemented without holding per-flow state in the forwarders, although it does require enough buffers to cope with the QoS strategy, as would be the case with any QoS mechanism.

IV. QoS IMPLEMENTATION IN THE PURSUIT ARCHITECTURE

A. Stateless QoS-based Forwarding

As previously explained, QoS differentiation is achieved by introducing “virtual links” in each forwarder. Implementation wise, each virtual link has its own queue and shaper. The shaper is responsible for enforcing the maximum bandwidth limit. All the queues that belong to the same link end up in a scheduler, which is responsible for enforcing priorities between them.

B. QoS Information in the Topology Manager

QoS information is mainly utilized by the TM function, since it is the one to select the most appropriate path through the network. However, the TM has to acquire the appropriate QoS information dynamically, which is viable with the use of Algorithmic Identifiers (AlgIds). An AlgId is automatically generated from a shared knowledge between the publisher and subscriber, e.g. by hashing a known name or existing identifier.

As mentioned in Section III-A, a hierarchy of information is maintained by the RV, in which each scope or item can have a number of QoS attributes associated with it. The TM receives a request from the RV to calculate a path between a publisher and a subscriber. The request includes the SId/RIId of the Information Item, from which the TM is able to calculate the AlgId of the associated QoS item by applying a hash function. For example, the QoS information of item /A/B/C is published under item /A/B/hash(C).

Now, the TM has the ability to subscribe to the QoS item and acquire the published requirements but a problem arises: The TM has already received the request for path formation while it currently has no knowledge about the associated QoS information. For this reason a short-term caching mechanism is employed where the TM stores the acquired QoS information. The steps the TM follows for constructing a QoS-based delivery path are:

- 1) Receive a path formation request from the RV
- 2) Check the QoS information cache for a valid entry:
 - If valid information is found for the item: use it for path calculation.
 - Else if valid information is not found for the item but QoS information is available for the parent scope: If the “delegate” property is set, use the parent’s QoS information for the item.
 - Else if no valid information is found:
 - Calculate the AlgId of the QoS item.
 - Subscribe to the QoS item (with a given timeout).
 - If QoS information is not published: Use the default/generic path calculation.
- 3) Calculate the correct FID and notify the appropriate publisher.

It should be noted that once the TM acquires the QoS information, and thus the information item’s priority, it has to select the links to be used in the calculation of the best path. In this step, the TM has to ensure that Low Priority (LP) items/

TABLE I
EXAMPLE WEIGHT MAP OF THREE VIRTUAL LINKS WITH PRIORITIES 0, 50 AND 90 BETWEEN NODES A AND B.

np / lp	AB_0	AB_{50}	AB_{90}
0	100	255	255
50	100	50	255
90	100	50	10

relations will avoid using High Priority (HP) links. This is achieved with the use of a weight map that efficiently (without multiple copies of the graph) slices the network and maps each relation to the best available QoS queues. A maximum of 100 different link priorities in the network is assumed. Taking for example three links between nodes A and B with priorities 0, 50 and 90 respectively, the corresponding weighting map can be seen in Table I. Each row represents a network priority and a set of edge weights to be used for Shortest Path (SP) calculation where each cell is calculated as:

$$weight = \begin{cases} 100 - lp, & \text{if } lp \leq np \\ 255, & \text{otherwise} \end{cases} \quad (1)$$

where lp is the current link priority (column) and np is the network priority (row). The TM selects the appropriate weight set by mapping the QoS information of the pub/sub relation to the closest, lower or equal available priority. Note that depending on the algorithm used, LP relations will prefer not to use HP links. Additionally, in the above example a HP relation may use LP links. This behavior can be adjusted by changing the first condition of Eq. 1 to equality.

On the other hand, if QoS information has not been published, the TM will prefer to use the default matching algorithm. This is preferred since extra matching delay is avoided. Additionally, when/if an item’s QoS information arrives at the TM, the TM will update the cache and optionally re-route the corresponding information item, now taking the QoS information into consideration. Note the delegate property of the QoS hierarchy in step 2 above, which allows multiple information items to share the same QoS information in a per-class like approach.

V. EVALUATION

Our model has been experimentally tested using the PURSUIT testbed for three classes of service: High Priority (HP): permitted to use all the link’s bandwidth, Medium Priority (MP): allowed to utilize up to 20% of the link’s bandwidth and, Low Priority (LP): allowed to utilize what is left of the link’s bandwidth. Experiments have shown that HP pub/sub relations can achieve an average bit rate of 130 Mbps with a mean inter-arrival time of 0.06 ms and mean jitter of 0.03ms.

Furthermore, we evaluate the application of our model over a realistic network scenario, modeled using the Janet backbone network topology [16], in terms of bandwidth saving, due to early traffic policing. This is achieved by the TM, which prohibits the load of a link’s class to increase beyond the assigned fraction of the link’s bandwidth. Thereby, avoiding traffic overflow at aggregation points and the waste of link bandwidth, occupied by the contending traffic. We, also,

evaluate the effect of early traffic policing on the drop rate of the offered load, expressed as a fraction of the overall offered load. This is achieved by comparing the drop rate of our model to that of IP DiffServ. Two capacity planning schemes have been considered, originally defined for MPLS networks: The Maximum Allocation Model (MAM) [17], and the Russian Doll Model (RDM) [18]. However, we do not employ TE objectives for path calculation to ensure fairness in comparison with IP DiffServ. We evaluate the trade-off between the MAM and RDM schemes in terms of fairness in sharing bandwidth among QoS classes. This is illustrated by comparing the load at the most congested link of a class with that of the other classes at the same link. In addition, we evaluate the message overhead introduced by the TM subscription to the QoS items and compare it with that of MPLS-TE.

Similar to our experiments, three classes of service have been defined with a per-class bandwidth allocation (for the MAM) of: 60% for HP traffic, 30% for MP traffic and, 10% for LP traffic. For our evaluation, we have designed an analytical model that expresses the total offered load as a fraction of the total network capacity. The offered load increases from 0.1 to 1 with steps of 0.1; at each step the network is tested for 100 experiments of randomized set of pub/sub relations. Message overhead has been averaged for all the experiments, where a HA-QoS scenario is assumed that allows an average of 10 relations to be associated with a single QoS scope.

A. Comparison with IP DiffServ

The results presented in Fig. 3 show the waste of bandwidth introduced by IP DiffServ, which is saved in our approach due to early traffic policing. For low load ratios (≤ 0.5), the saving in the HP class is lower than that of MP and LP classes. This is due to the higher bandwidth allowance of the HP class, which results in, generally, lower blocking rate as shown in Fig. 4. However, as the load increases, congestion on the critical links in the network also increases. Consequently, the dropping rate of HP traffic at the critical links increases and therefore results in a higher bandwidth waste at the critical as well as the prior links (i.e. from the source to the dropping point). As a result, a higher bandwidth saving is exhibited in the HP class compared to MP and LP classes, for load ratios higher than 0.5, with an approximate average bandwidth saving of 10%. The results exhibited by the MP and LP classes are, in general, lower with respect to the overall network capacity due to the lower fraction of bandwidth allowance for each class. As a result, traffic dropping occurs at closer links to the source, where the number of links traversed so far is relatively small and therefore, bandwidth waste is smaller than that shown in the HP and thus the benefit is present but not as obvious.

The results presented in Fig. 4 show a significant reduction in the drop ratio of the PURSUIT models (MAM and RDM) compared to IP DiffServ. This is due to the early traffic policing, which allows the preserved resources to be utilized for further pub/sub relations that can be admitted to the network without causing a traffic contention. It can be noticed that the best results are achieved by the PURSUIT-RDM




Fig. 3. Bandwidth saving compared to IP DiffServ where the saving is expressed as a fraction of the network total capacity




Fig. 4. Reduction in the drop ratio of pub/sub relations compared to IP DiffServ where the drop ratio is expressed as a fraction of the total offered load per class

model. This is due to the flexibility in allowing all the classes to utilize all of the link's bandwidth. As a result, the ability of the network to admit requests of lower classes over paths that are underutilized by higher classes increases. However, the flexibility in bandwidth allocation offered by the PURSUIT-RDM comes at the expense of allowing the higher classes to opt out all the lower classes. This is illustrated by the results presented in Fig. 5 and Fig. 6, which show the load on the most congest link of each class and the load on the other classes of the same link respectively. Although, in general, the PURSUIT-RDM achieves a better dropping rate than the PURSUIT-MAM, the latter's constraint on the capacity per class assignment ensures a fairer sharing of the link's capacity across all classes, and thereby avoids starvation scenarios of lower classes on congested links.

B. Comparison with MPLS-TE

With regards to message overhead, it is highly dependent on the QoS strategy and the number of pub/sub relations associated with a single strategy. However, within our analytical model and for one direction, an average of 77 subscription messages are issued by the TM to obtain the QoS information of an average of 800 pub/sub relations. Contrarily, an average of 2277 path messages are sent by MPLS-TE to admit the same set of flows. The results indicate the superiority of our model in terms of scalability compared to MPLS-TE.




Fig. 5. Utilization ratio per class of the most congest link where utilization ratio is expressed as a fraction of the total link capacity




Fig. 6. Utilization ratios of complementary class of the most congest link where utilization ratio is expressed as a fraction of the total link capacity

VI. CONCLUSION

QoS provisioning is an essential challenge for future Internet architectures, such as information-centric networks. In this work, we have presented a novel QoS model for the PURSUIT ICN architecture that achieves native, information-centric, QoS provisioning. We have shown the flexibility of the proposed model in supporting different QoS scenarios over the same pub/sub communication framework. Moreover, we have described the implementation of our model in the PURSUIT testbed and presented the experimental results of example pub/sub relations. We have evaluated the benefits of early traffic policing in our model, for two capacity planning schemes (MAM and RDM), and the fairness in bandwidth allocation among the classes, when applied to a realistic network. Modeling results indicate a significant bandwidth saving and drop rate reduction when compared to the IP DiffServ model. The results also indicate a better fairness achieved by MAM compared to RDM. Finally, we have illustrated the advantages of our model over MPLS-TE, in terms of scalability, by comparing the message overhead introduced by the two models in our experiments.

ACKNOWLEDGMENT

The work of the authors was funded by the EU framework program under contract no. ICT-2010-257217. The design considerations in this paper are the result of intensive discussions

among the researchers in the PURSUIT project and beyond.

REFERENCES

- [1] M. L. Mueller and H. Asghari, "Deep packet inspection and bandwidth management: Battles over bittorrent in canada and the united states," *Telecommunications Policy*, vol. 36, no. 6, pp. 462 – 475, 2012, global Internet Governance Research and Public Policy Challenges for the Next Decade. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0308596112000523>
- [2] J. W. Evans and C. Filsfil, *Deploying IP and MPLS QoS for Multiservice Networks: Theory & Practice*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2007.
- [3] A. Abella, V. Friderikos, and H. Aghvami, "Differentiated services versus over-provisioned best-effort for pure-IP mobile networks," in *Mobile and Wireless Communications Network, 2002. 4th International Workshop on*. Ieee, 2002, pp. 450–457. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1045806>
- [4] U. Fiedler, P. Huang, and B. Plattner, "Over-provisioning or Differentiated Services - A Case Study on integrating services over IP," *Tech. Rep. 1*, Oct 2000.
- [5] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao, "Overview and principles of internet traffic engineering," United States, 2002. [Online]. Available: www.ietf.org/rfc/rfc3272.txt
- [6] D. Zhang and D. Ionescu, "Qos performance analysis in deployment of diffserv-aware mpls traffic engineering," in *Proceedings of the Eighth ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing - Volume 03*, ser. SNPD '07. IEEE Computer Society, 2007, pp. 963–967.
- [7] V. Jacobson, D. K. Smetters, J. D. Thornton, M. Plass, N. Briggs, and R. Braynard, "Networking named content," *Commun. ACM*, vol. 55, no. 1, pp. 117–124, Jan. 2012.
- [8] D. Trossen, M. Särelä, and K. Sollins, "Arguments for an information-centric internetworking architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 40, pp. 26–33, 2010.
- [9] N. Fotiou, P. Nikander, D. Trossen, and G. C. Polyzos, "Developing Information Networking Further: From PSIRP to PURSUIT," in *International ICST Conference on Broadband Communications, Networks, and Systems (BROADNETS), 2010 (invited paper)*, Oct 2010.
- [10] G. Pavlou, N. Wang, W. K. Chai, and I. Psaras, "Internet-scale content mediation in information-centric networks," *Annales des Télécommunications*, vol. 68, no. 3-4, pp. 167–177, 2013.
- [11] T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 181–192, Aug. 2007.
- [12] P. Jokela, A. Zahemszky, C. Rothenberg, S. Arianfar, and P. Nikander, "Lipsin: line speed publish/subscribe inter-networking," *SIGCOMM Comput. Commun. Rev.*, vol. 39, pp. 195–206, 2009.
- [13] G. Xylomenos, C. Ververidis, V. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. Katsaros, and G. Polyzos, "A survey of information-centric networking research," *Communications Surveys Tutorials, IEEE*, accepted for publication in 2013, available as preprint from <http://dx.doi.org/10.1109/SURV.2013.070813.00063>.
- [14] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *Communications Magazine, IEEE*, vol. 50, no. 7, pp. 26–36, July 2012.
- [15] W. K. Chai, I. Psaras, M. Charalambides, D. Tuncer, G. Pavlou, D. Flórez, F. J. Hernández Romero, A. Beben, J. M. Batalla, J. Sliwinski, P. Wisniewski, W. Burakowski, N. Wang, G. Kamel, S. Spirou, and M. Georgiades, "D3.2 content mediator architecture for content-aware networks european seventh framework project fp7-2010-ict-248784-strep final specification of mechanisms, protocols and algorithms for the content mediation system," pp. 1–100, 2012.
- [16] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The internet topology zoo," *IEEE Journal on Selected Areas in Communications*, vol. 29, pp. 1765–1775, 2011.
- [17] F. L. Faucheur and W. Lai, "Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering," Internet Engineering Task Force, June 2005. [Online]. Available: <http://www.ietf.org/rfc/rfc4125.txt>
- [18] F. L. Faucheur, "Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering," Internet Engineering Task Force, June 2005. [Online]. Available: <http://www.ietf.org/rfc/rfc4127.txt>