QOS-DRIVEN MULTIPATH ROUTING FOR ON-DEMAND VIDEO STREAMING IN A PUBLISH-SUBSCRIBE INTERNET

Yannis Thomas, Pantelis A. Frangoudis§, George C. Polyzos**

*Department of Informatics, Athens University of Economics and Business, Greece
§IRISA/INRIA Rennes-Bretagne Atlantique, France
thomasi@aueb.gr, pantelis.frangoudis@inria.fr, polyzos@aueb.gr

ABSTRACT

Aiming to improve the performance of on-demand video streaming services in an Information-Centric Network, we propose a mechanism for selecting multiple delivery paths, satisfying bandwidth, error rate and number-of-paths constraints. Our scheme is developed in the context of the Publish-Subscribe Internet architecture and is shown to outperform state-of-the-art multi-constrained multipath selection mechanisms by up to 7%, and single-path or single-constrained multipath selection schemes by up to 17%, in terms of feasible path discovery, while at the same time improving on bandwidth aggregation. Also, it is suitable for supporting resource-demanding high-definition scalable video streaming services, offering Quality-of-Experience gains.

1. INTRODUCTION

Information-Centric Network (ICN) architectures [1] have emerged as a response to the shift in traditional Internet usage patterns: The host-to-host communication paradigm is being reconsidered in favor of a content-centric one, where the interest lies in the information itself, rather than its actual origin. This shift is typically manifested in video services, where content is distributed by means of Content Delivery Networks (CDNs), where the actual host(s) serving the content is of little importance to its consumers.

With video traffic dominating [2] and the demands for HD content anywhere, anytime, and on any device growing [3], more efficient delivery techniques and the necessary network- and application-level support are being called for. At the network level, the exploitation of multiple delivery paths simultaneously is considered, offering bandwidth aggregation and resilience advantages. At the application level, advanced encoding techniques, such as Scalable Video Coding (H.264/SVC) [4], often coupled with adaptive streaming, can more efficiently cope with the dynamics of network conditions and resource demand.

ICN architectures, and the Publish Subscribe Internet (PSI) architecture [5] in particular, in the context of which we position our work, naturally support multipath delivery. In this paper, we make the following contributions: (i) We show where and how this functionality fits in the overall PSI architecture, (ii) propose an algorithm for multiple path selection which caters for multiple performance criteria, showing it to outperform state-of-the-art path selection schemes, and (iii) demonstrate how our multipath scheme can be used by a scalable video delivery application, quantifying its performance advantages in QoE terms.

2. BACKGROUND

2.1. Publish Subscribe Internet

The Publish Subscribe Internet (PSI) is an ICN architecture based on the *pub-sub* paradigm. Each *information item* is assigned a statistically unique identifier for addressing it, regardless of its location or owner. Content requests are addressed to the network, contrary to IP's end-to-end interaction, maintaining the loose coupling of pub-sub: A *publisher* (content provider) issues a publication about an information item and the network stores this information. When a *subscriber* (content consumer) issues a subscription to this content, the network matches it with the prior publication(s), locates it, and undertakes its delivery to the requesting host. This is performed via three clearly separated core functionalities, i.e., Rendezvous, Topology and Forwarding.

Topology functionality involves the discovery of dissemination routes between a publisher and multiple subscribers[1] and Forwarding functionality is responsible for content delivery via source routing. Path discovery is executed by the Topology Manager (TM), a logical entity which receives requests from the network's Rendezvous elements, formulates the appropriate transmission paths and sends them to the end users for direct communication, aiming to satisfy communi-

This work has been co-financed by the European Union (European Social Fund-ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) Research Funding Program: Aristeia II/I-CAN, and by the French Ministry of Economy, Industry and Digital Affairs through the Celtic-Plus project QuEEN (Project ID: CP8004)

[1] PSI natively operates on multicast, but unicast is also supported.

cation needs by being aware of the complete network topology (including link propagation delays, capacities and error rates). A request to the TM also carries a "strategy" flag, which defines the preferred dissemination pattern, such as multipath, multisource, unicast, etc. The TM is where traffic engineering schemes that improve network operation and adapt to individual service requirements can be applied, and this is the context for our work on multipath QoS routing.

The centralized nature of the Topology functionality has raised concerns about the feasibility of PSI. The TM service computes the data paths of roughly all network connections, which is questioned by network engineers who prefer pushing computational costs towards the edge of the network. Nevertheless, Alzahrani et al. [6] validate that an intra-domain TM solution is feasible and affordable by a reasonable number of TM instances with precomputed paths. Moreover, since Internet traffic is dominated by video streaming, deploying dedicated TM instances that discover paths only for streaming-related services is an appealing solution.

2.2. Path diversity for video delivery in IP networks

The suitability of multipath for multistream video coding is well known [7, 8]. To overcome IP-related limitations on exploiting path diversity, three methods are mainly assumed: multisource, overlay networks and MPLS-based routing. Multisource video streaming is widely used by CDNs for providing low latency and scalability. Although outside the scope of this paper, our multipath design could further enhance multisource gains.

In application-based overlay networks, a user sends each stream to a different *relay* node as an intermediate communication step. The number and location of relays highly affects the performance of this solution, wrt. resilience and latency. For error rate minimization, a relay must be placed at every branching node, making this method too expensive.

Finally, MPLS based routing allows source routing in IP networks via label stacking. MPLS comes with apparent costs (additional infrastructure requirements, increased traffic footprint, stateful routing), that limit its application to backbone networks, in turn limiting the gains of exploiting topological richness. Additionally, MPLS rationale is not in line with multipath video streaming semantics. MPLS assigns priorities to services and all packets of the same service follow the same route. Therefore in order to allow multipath video streaming, the characterization level must be lowered to bitstreams. Furthermore, using ingress MPLS routers to split the bistream among different paths results in static substream-path allocation, since the MLPS routers do not support packet scheduling algorithms [9, 10] for video streaming.

On the other hand, source routing is native in the entire network in PSI, without requiring intelligent routing components, thus maximising path diversity. Moreover, bitstream scheduling among paths in PSI lies on the application, instead of the ingress MPLS router, allowing agile resource utilization and efficient packet selection and scheduling.

2.3. QoS routing and path formation

Path formation algorithms for QoS routing are classified in three categories wrt. their complexity, and thus their feasibility [11]. *Single-Constraint Optimization Path* (SCOP) algorithms optimize for a single metric, commonly latency or hop-count, and run in polynomial time. *Multi-Constrained Optimization Path* (MCOP) algorithms discover optimal paths under numerous constraints and objectives, but typically no polynomial-time solutions exist. Finally, *Multi-Constrained Path* (MCP) algorithms select paths that satisfy multiple constraints without optimization; any feasible path is acceptable.

Wang and Crowcroft [12] exploit the fact that bandwidth is a concave metric and propose a polynomial-time algorithm for multi-constrained routing. Their scheme initially prunes all network links with capacity lower than the bandwidth requirement and, then, the Dijkstra algorithm discovers the shortest available path on the modified graph. This algorithm balances between the MCOP and MCP categories, optimizing delay but accepting any path with bandwidth over the constraint. We are aware of only two multi-constrained *multipath* QoS routing algorithms, MADSWIP [13] and DIMCRA [14]. MADSWIP (MCOP) computes maximally-disjoint widest paths, with Dijkstra-like complexity. It supports multi-constrained optimization via the lexicographical comparison of links. First, it minimizes shared links, then it maximizes bandwidth and, finally, it minimizes delay. Due to this "hierarchical optimization," MADSWIP fails to select feasible paths when the delay of the maximally-disjoint paths violates the delay constraint. DIMCRA (MCP) attempts to discover two disjoint paths wrt. multiple constrains. It exploits tree pruning and applies SAMCRA, an MCP algorithm for unicast routing, so as to form the two optimal, loop-free, disjoint paths.

Finally, MMSPEED [15] is a multi-constrained multipath routing scheme for wireless sensor networks. It optimizes latency and reliability by placing logic at the forwarding level; computationally-enhanced network routers decide the next destination of each packet based on localized knowledge. This pattern can be facilitated in PSI, too, either by PSI routers or by TMs, provided that TMs are updated with the required intelligence. MMSPEED is more of a real-time forwarding method than a path selection algorithm, therefore we currently omit further investigation.

3. A MULTI-CONSTRAINED MULTIPATH SELECTION ALGORITHM FOR PSI

We introduce a Bandwidth-, Error-rate- and number-of-Paths-constrained MultiPath (BEPMP) selection algorithm, a polynomial time MCP scheme which, to the best of our knowledge, is the first to support QoS routing operating under these three constraints. It addresses the needs for high bandwidth and low loss rates of on-demand video streaming services and supports constraints on the maximum number of available paths, motivated by (i) the need for lower path discovery algorithmic complexity and lower path management complexity from an application perspective, and (ii) the fact that a large number of paths does not necessarily improve path diversity, which highly depends on the nature of the topology [16].

BEPMP operates in three steps, each one ensuring the satisfaction of a single constraint. Initially, it selects the paths that satisfy the error rate constraint by running Yen's k-shortest paths algorithm [17] with the error rate as the metric to minimize. Yen's algorithm returns a sorted list of source-destination paths. BEPMP keeps only the first n of these paths which do not violate the error rate constraint. Then, it selects a set of m out of the n paths that jointly satisfy the bandwidth constraint, using a variation [18] of the Ford-Fulkerson (FF) maximum flow algorithm. In particular, it runs its "widest augmenting path" (WAP) version on the topology graph composed only of the n paths returned by Yen's algorithm, selects the widest path (i.e., the one that admits the most flow) at each step, and stops either (i) when the network flow calculated covers the bandwidth constraint, or (ii) after the selection of the m widest paths; if the aggregate bandwidth of these m paths is lower than the constraint, the algorithm decides that no feasible solution exists.

Data: Graph G, constraints BW, $ERROR$, m
Result: A set of at most m paths with
$cumulative\ capacity > BW$ and
$maximum\ error\ rate < ERROR$
begin
    Get k source-destination paths that minimize error rate;
    Drop paths that do not satisfy $ERROR$ constraint;
    while *all constraints can be met* **do**
        Select the widest path of the remaining;
        Estimate cumulative bandwidth;
        Update residual graph;
    end
end

Algorithm 1: The BEPMP algorithm

BEPMP has polynomial complexity, involving the sequential execution of two polynomial algorithms. Yen's k-shortest paths algorithm has an $O(N^2)$ worst-case complexity on the number of nodes, while the WAP maximum flow algorithm runs in $O(E^2\ logE\ logf^*)$ time [18], where E is the number of edges (links) and $f*$ is the maximum flow (aggregate bandwidth) value. BEPMP finds a feasible path if this exists (*correctness*) if Yen's algorithm is configured with k high enough so that BEPMP searches among all error-rate-feasible paths. If k is very small, it is likely that a path which can satisfy both the error rate and bandwidth constraints exists, but BEPMP fails to discover it because it was not one of the top-k paths. Adaptively setting k based on the characteristics of the topology is a topic we defer for future work.

Metric	Average	Standard Deviation
#Nodes	29.6	14.3
#A-Nodes	11.1	10.5
#Edges / #Nodes	1.45	0.52
Avg. link BW	2.1 Gbps	1.98 Gbps

Table 1: Statistics of the topologies used.

4. EVALUATION

4.1. Implementation

We have built BEPMP, DIMCRA, Dijkstra, Wang's and Yen's algorithms into the PSI TM. Our TM implements path pre-computation [6] using any of the above algorithms, computing all available dissemination paths among all pairs of access nodes (A-Nodes)[2] of the topology. Note that we constrain BEPMP's and Yen's operation to a maximum of two paths, in order to be comparable with DIMCRA. We also allow BEPMP and DIMCRA to terminate when the first path meets all constraints, thus returning a single path.

4.2. Topologies

We experimented with a variety of topologies investigating the effects of network size and density. We randomly selected 20 AS-level topologies from topology-zoo[3] that include link capacities; the topology characteristics are illustrated in Table 1. We assign link error rates following the Zipf distribution, resulting in many links with lower error probability (approximately 0.01%) and a few with a substantial one (up to 1%). The latter could correspond to, e.g., lossy wireless links. We further introduce random network load via uniformly-distributed artificial competitive flows: We select x paths from a random source node to a random A-Node, and we remove from each path link $y\%$ of the capacity of the path's first link, where $x = \#Nodes \times 10$ and $y = 0.5$. Each presented result (Figures 1, 2) is the mean across all topologies, for 100 iterations for each topology, each with varying error rates and network load.

4.3. Service ratio

We define the number of feasible path computations divided by the total number of path computations as the *service ratio*. This metric illustrates an algorithm's efficiency in pro-

[2] A-Nodes have a link degree equal to 1. If a topology includes less than four A-Nodes, we insert random ones.

[3] http://www.topology-zoo.org/

viding routes for successful communication. In our experiments (Figure 1), we first run the Dijkstra algorithm to find the "shortest" (minimum error rate) path between *all pairs of A-nodes*, and we estimate the average available connection bandwidth and error rate of each topology. These values are used as baseline constraints. We then deploy Yen's, Wang's, DIMCRA and BEPMP algorithms to find the paths that accommodate certain service requirements among all A-nodes. Since Dijktra and Yen do not support a second constraint, we apply the different requirements on the computed paths a posteriori.

We first examine the performance of the algorithms as the bandwidth requirement increases from 50% to 150% of the baseline requirement and the error constraint remains at 100%. BEPMP scores better than any other algorithm. Yen's algorithm scores the lowest, as it ignores the bandwidth requirement and always computes a second path (even when not needed), potentially increasing the overall error rate. Finally, the performance of all algorithms degrades as the capacity constraint raises, but their ranking does not change. Unicast algorithms suffer 15% greater reduction than multipath ones, validating the gains of multipath bandwidth aggregation.

Then, we increase the error rate constraint from 50% to 150% of the baseline requirement, keeping the bandwidth requirement at 100%. The algorithms rank the same as in the previous scenario; loosening the constraint affects all algorithms roughly the same. BEPMP responds slightly better to high error rate requirements than DIMCRA, as their performance difference increases from 3% to 7% in the 50% and 150% error rate constraint cases respectively. This is also expected given that BEPMP allows overlapping paths, while DIMCRA forms disjoint paths only, thus potentially increasing the service error rate. Overall, BEPMP achieves better service ratio in all investigated constraint setups (up to 7% better service ratio than DIMCRA, and up to 17% better than any of the other algorithms).

4.4. Second path utilization

We now focus on the utilization of the second path. We omit Yen's algorithm because it always uses a second path, contrary to DIMCRA and BEPMP that terminate when a single path satisfies both constraints. Again, we vary one constraint from 50% to 150% of the value of the baseline constraint, while keeping the second constraint unchanged (Figure 2).

We confirm three important arguments in favor of multipath transmissions. First, multipath communications serve throughput-intensive applications better. The more bandwidth a service demands, the more bottleneck-disjoint paths will be exploited; DIMCRA increases second path utilization roughly 17% as the bandwidth constraint increases. Furthermore, using multiple paths may increase the error rate of the service. In our experiments both algorithms increase the second path utilization roughly by 16% when the application per-

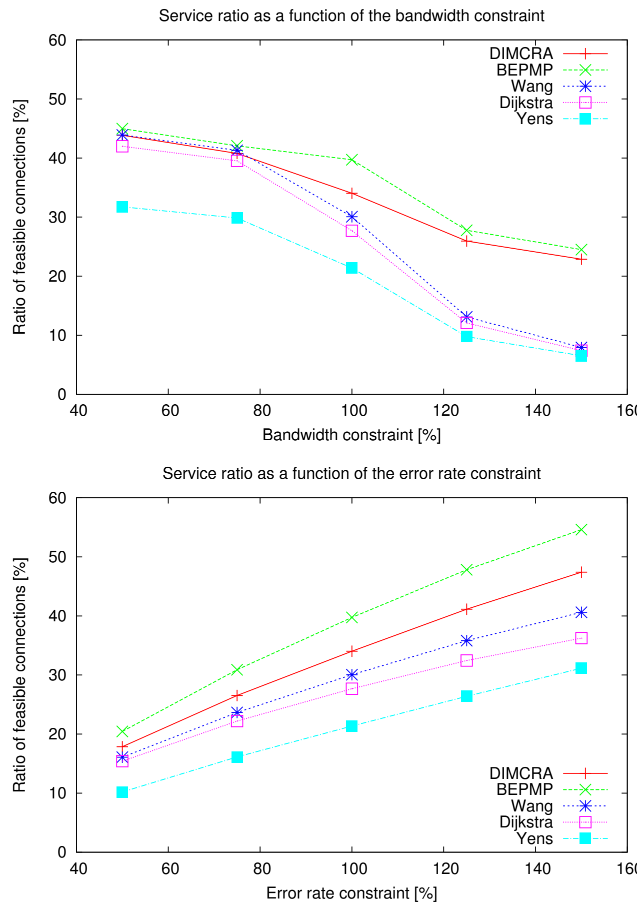

Fig. 1: Service ratio of the different algorithms under various bandwidth and error rate requirements.

mits higher error rates. While selecting a second path does not necessarily increase the overall error rate, our results depict a proportional relationship of the two quantities in practice. Finally, we verify that allowing only disjoint paths penalizes the applicability of multi-flow connections. DIMCRA uses a second path roughly 10% less than BEPMP, which allows overlapping paths. Multipath transmission must avoid the shared bottlenecks and not just *any* shared link. There are wide links that can be used by both sub-streams without causing congestion. Exploiting such shared links improves the performance of BEPMP compared to DIMCRA.

4.5. QoE benefits for an SVC video streaming application

4.5.1. Application settings

As a use case, we study the interplay between our path selection algorithm and a video streaming application. In particular, we assume a video-on-demand streaming service for the delivery of high-definition content. Scalable video technologies are used, so that a layered representation of the content is possible, where a lower-bitrate base layer is necessary for

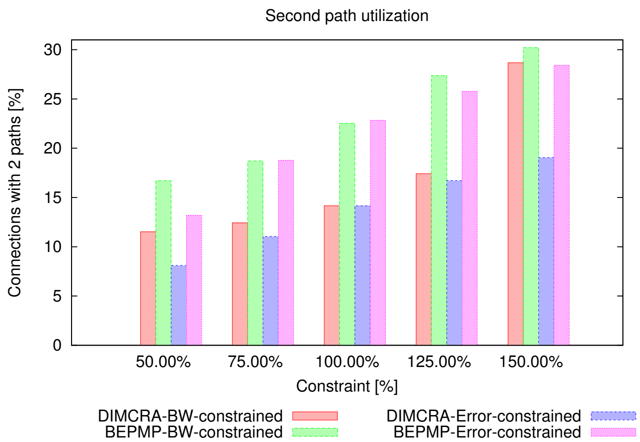

Fig. 2: Ratio of feasible connections that utilize two paths.

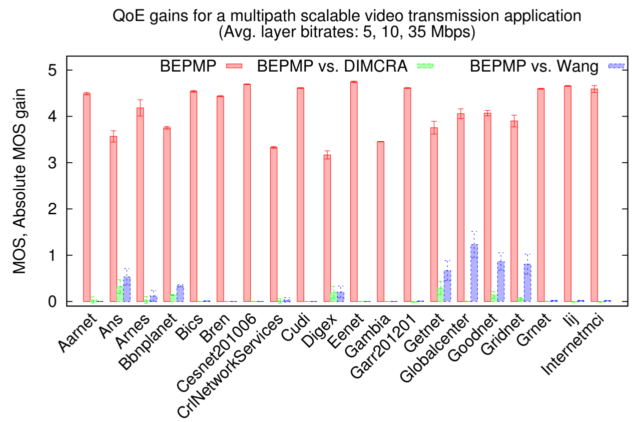

Fig. 3: Comparison of the QoE achieved by an SVC HD video delivery application for different path selection algorithms. We present the estimated MOS for BEPMP, and the difference in MOS compared to the case when the DIMCRA and Wang algorithms are used, for 20 different topologies.

decoding, and a number of enhancement layers are available for improved video quality (but also increased bitrate).

We assume that the H.264/SVC [4] encoding technology is used, and three video layers are provided. The video streaming application distributes the video over the paths returned by the path selection algorithm so that user experience is maximized. This is compatible with the PSI design: Using a "slow-path" rendezvous mechanism, the (up to) two paths and their specifications (end-to-end packet loss rate, available bandwidth), are available to the application. Then, the latter can select how to stream data over the available paths on a per information item (or even on a per packet) basis, aiming to maximize user experience.

The application logic involves the following steps: (i) Calculate all possible assignments of video layers to the available paths. This number is typically very small (e.g., if 3 layers and 2 paths are available, 8 potential assignments are evaluated). (ii) For each assignment, check if the bandwidth requirements of the video are covered. If not, drop enhancement layers until there is enough capacity for streaming. (iii) Calculate the QoE expected to be achieved under the specific assignment using a QoE model for SVC video (see Section 4.5.2). (iv) Select the assignment which maximizes the estimated QoE.

4.5.2. QoE assessment

To quantify QoE, we apply the Pseudo-Subjective Quality Assessment (PSQA) [19] approach, which involves training a Random Neural Network (RNN) on subjective tests under controlled conditions, where a set of parameters affecting quality is monitored and the ratings of users are recorded. The trained RNN classifier can then be applied in real time and output the expected mean opinion score (MOS) on the 1-5 (poor-excellent) scale for specific values of the input parameters. In this work, we use an instance of PSQA trained for QoE estimation of scalable video [20], where fixed values for the frame rate and the spatial resolution are assumed, and the video contains three SNR-scalable layers. The input parameters are the frequency of Instantaneous Decoder Refresh (IDR) video frames[4] and the loss rate per layer.

The application uses this tool to select the appropriate path for each layer, assuming a fixed IDR frequency value (fIDR = 30) and using the loss information per path as input to the model. Note that if a layer cannot be accommodated by a path and is dropped, the loss rate for this layer, but also for higher layers which depend on it for decoding, are set to 100%.

4.5.3. Performance benefits

We apply the layer-path assignment mechanism described in Section 4.5.1 for the paths returned by three different algorithms: (i) DIMCRA, (ii) Wang (single-path), and (iii) BEPMP. We present a comparison of the three schemes across 20 different topologies, for 100 iterations for each topology, where we vary the link loss rates and the available link capacities as a result of parallel flows. Error rate and bandwidth constraints are set to 100% of the baseline case (see Section 4.3). For each iteration, we calculate the mean QoE across source-destination pairs for each algorithm. To experiment with realistic HD video bandwidth requirements, we used data available from an ISO/IEC SVC verification test report [21]. In particular, we assume the transmission of a three-layer 1080p SVC video, where the bitrates of the base layer and the two enhancement layers are $b_0 = 5$Mbps, $b_1 = 10$Mbps, and

[4] IDR frames are encoded without reference to any other frames, as is the case, e.g., for B or P frames in an MPEG video. Thus, they can also be decoded independently, and the higher their frequency, the more resilient the video is to losses, at the expense of higher bitrate.

$b_2 = 35$Mbps, respectively. Figure 3 presents the results of our study. Each reported value is the mean of all iterations presented with 95% confidence intervals. Our path selection algorithm outperforms the other two schemes in most of the examined topologies, while at the same time achieving higher service ratio.

5. CONCLUSION AND FUTURE WORK

We presented BEPMP, a multipath selection strategy for the PSI ICN architecture, which simultaneously considers bandwidth, error rate and number-of-paths constraints. Our polynomial-time algorithm admits centralized implementation and is feasible for supporting QoS routing at the intra-domain level. Experimenting with diverse network topologies, we have shown it to outperform state-of-the-art multi-path selection schemes. We have further shown that QoE-aware scalable video applications for on-demand multipath video streaming over PSI can be built on top of BEPMP, and quantified the performance improvements it brings in user experience terms. This work will serve as a basis for our future research on multiple fronts, including BEPMP algorithm refinements, extensions towards multi-source video delivery, particularly addressing environments with dynamic video source presence, and, importantly, a large-scale experimental evaluation on a testbed like Planetlab.

6. REFERENCES

[1] G. Xylomenos, C.N. Ververidis, V.A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K.V. Katsaros, and G.C. Polyzos, "A survey of information-centric networking research," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 2, pp. 1024–1049, 2014.

[2] Cisco, "Cisco Visual Networking Index: Forecast and Methodology, 2013-2018," White Paper, June 2014.

[3] Frost & Sullivan, "Global Pay TV Middleware Market," Market Research, November 2014.

[4] H. Schwarz, D. Marpe, and T. Wiegand, "Overview of the Scalable Video Coding Extension of the H.264/AVC Standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 17, no. 9, pp. 1103–1120, Sept. 2007.

[5] "PURSUIT: Publish-Subscribe Internet Technology," `http://www.fp7-pursuit.eu`.

[6] B.A. Alzahrani, M.J. Reed, J. Riihijärvi, and V.G. Vassilakis, "Scalability of information centric networking using mediated topology management," *J. Netw. Comput. Appl.*, 2014.

[7] J.G. Apostolopoulos and M.D. Trott, "Path diversity for enhanced media streaming," *IEEE Commun. Mag.*, vol. 42, no. 8, pp. 80–87, 2004.

[8] A.C. Begen, Y. Altunbasak, O. Ergun, and M.H. Ammar, "Multi-path selection for multiple description video streaming over overlay networks," *Signal Process.: Image Commun.*, vol. 20, no. 1, pp. 39–60, 2005.

[9] D. Jurca and P. Frossard, "Video packet selection and scheduling for multipath streaming," *IEEE Trans. Multimedia*, vol. 9, no. 3, pp. 629–641, 2007.

[10] V. Singh, S. Ahsan, and J. Ott, "Mprtp: multipath considerations for real-time media," in *Proc. ACM MMSys*, 2013, pp. 190–201.

[11] F. Kuipers, P. Van Mieghem, T. Korkmaz, and M. Krunz, "An overview of constraint-based path selection algorithms for qos routing," *IEEE Commun. Mag.*, vol. 40, no. 12, 2002.

[12] Z. Wang and J. Crowcroft, "Quality-of-service routing for supporting multimedia applications," *IEEE JSAC*, vol. 14, no. 7, pp. 1228–1234, 1996.

[13] N. Taft-Plotkin, B. Bellur, and R. Ogier, "Quality-of-service routing using maximally disjoint paths," in *Proc. IEEE IWQoS*, 1999, pp. 119–128.

[14] Y. Guo, F. Kuipers, and P. Van Mieghem, "Link-disjoint paths for reliable qos routing," *Int. J. Commun. Syst.*, vol. 16, no. 9, pp. 779–798, 2003.

[15] E. Felemban, C.-G. Lee, E. Ekici, R. Boder, and S. Vural, "Probabilistic qos guarantee in reliability and timeliness domains in wireless sensor networks," in *Proc. IEEE INFOCOM*, 2005, vol. 4, pp. 2646–2657.

[16] J.P. Rohrer, A. Jabbar, and J.P.G. Sterbenz, "Path diversification for future internet end-to-end resilience and survivability," *Telecommun. Syst.*, vol. 56, no. 1, pp. 49–67, 2014.

[17] J.Y. Yen, "Finding the k Shortest Loopless Paths in a Network," *Manage. Sci.*, vol. 17, no. 11, pp. 712–716, July 1971.

[18] J. Edmonds and R.M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," *J. ACM*, vol. 19, no. 2, pp. 248–264, April 1972.

[19] G. Rubino, "Quantifying the Quality of Audio and Video Transmissions over the Internet: The PSQA Approach," in *Communication Networks & Computer Systems*, J.A. Barria, Ed. Imperial College Press, 2005.

[20] K.D. Singh, A. Ksentini, and B. Marienval, "Quality of Experience Measurement Tool for SVC Video Coding," in *Proc. IEEE ICC*, 2011.

[21] ISO/IET JTC 1/SC29/WG 11, "SVC Verification Test Report," Tech. Rep. N9577, 2007.