

Federated and Autonomic Management of Multimedia Services

Jeroen Famaey and Filip De Turck

Department of Information Technology, Ghent University – iMinds

Abstract—Over the years, the Internet has significantly evolved in size and complexity. Additionally, the modern multimedia services it offers have considerably more stringent Quality of Service (QoS) requirements than traditional static services. These factors contribute to the ever-increasing complexity and cost to manage the Internet and its services. In the dissertation, a novel network management architecture is proposed to overcome these problems. It supports QoS-guarantees of multimedia services across the Internet, by setting up end-to-end network federations. A network federation is defined as a persistent cross-organizational agreement that enables the cooperating networks to share capabilities. Additionally, the architecture incorporates aspects from autonomic network management to tackle the ever-growing management complexity of modern communications networks. Specifically, a hierarchical approach is presented, which guarantees scalable collaboration of huge amounts of self-governing autonomic management components.

I. INTRODUCTION

Since its inception, the Internet has grown from a small packet-switched computer network, known as ARPANET, into a global communications substrate, consisting of over 40,000 interconnected networks. Additionally, it has continuously evolved towards richer and more demanding services. The traditional static services, such as email, Usenet and the World Wide Web (WWW), can easily handle network hiccups (e.g., increased delay and jitter or reduced throughput). In contrast, the quality of modern multimedia services quickly deteriorates under suboptimal network conditions. These two evolutions have significantly contributed to the Internet's increased operational and management complexity.

Despite this metamorphosis, the Internet's ground rules and underlying architectural concepts have remained largely the same over the years. The networks that make up the global Internet still stand on their own, with little to no interaction or collaboration among them. Moreover, transmission and routing of data is still packet- and best-effort-based. This long-term stagnation has been the source of several recent problems. First, the ever-increasing size and complexity of the Internet has caused a shift in the costs associated with maintaining it. Where originally the majority of the costs were associated with the hard- and software infrastructure itself, this has shifted towards the cost of managing and supporting it [1]. Moreover, due to the Internet's best-effort nature and the lack of cooperation among interconnected networks, the stringent end-to-end requirements of modern multimedia services cannot be guaranteed [2].

To address these concerns, there is a need for novel management approaches. Specifically, *network federations* [3] have been advanced as a method for satisfying the end-to-end quality requirements of novel multimedia services, while

the *autonomic network management* [4] paradigm has been proposed to encompass the ever-increasing size and complexity of communications networks. Network federations are defined as persistent cross-organizational agreements that enable the cooperating networks to share capabilities in a controlled way. These capabilities might range from the reservation of bandwidth-guaranteed paths to control over specific device configurations. The autonomic network management paradigm aims to reduce the management complexity of large-scale communications networks for the human network operator. This is achieved by introducing self-governing capabilities into the network, which enable it to manage itself within the boundaries of the management policies specified by the operators.

In our work, a novel Future Internet architecture is presented that incorporates aspects of both federated and autonomic network management. Specifically, the contributions of our work are twofold. First, a novel methodology for the federated delivery of Internet-based multimedia services with stringent QoS requirements is presented. Second, several architectural components are introduced to facilitate the scalable collaboration between distributed self-governing management components in federated autonomic network management frameworks.

The remainder of this paper is structured as follows. Section II gives a high-level overview of the research contributions of the Ph.D. and offers insights on how they relate to each other. Additionally, it provides a list of publications and collaborations resulting from the work. Subsequently, Section III describes the first contribution, on federated delivery of multimedia services, in more detail. The second contribution, on scalable cooperation between autonomic management components, is discussed in Section IV. Finally, Section V lists the major conclusions of the thesis.

II. OVERVIEW OF RESEARCH CONTRIBUTIONS

In our work, we aimed to incorporate and combine aspects from the autonomic and federated network management paradigms in the management of the Future Internet and its services. The ultimate goal is to improve the Internet's ability to cope with its ever-increasing size and complexity, as well as the increasingly stringent QoS requirements of novel multimedia services. This translates into several contributions that are depicted in Figure 1.

The figure depicts an end-to-end view of the Internet, consisting of multiple independently managed network domains involved in our envisioned federated multimedia service delivery architecture. On top of the Internet's infrastructure




Fig. 1. A graphical depiction of the contributions of the thesis, positioned within an end-to-end view of the Internet

and services is a management layer, consisting of a huge number of self-governing autonomic management components, called Autonomic Elements (AEs). Through the Semantic Communications Bus (SCB), AEs communicate and cooperate both within and across domain boundaries, in order to achieve QoS-guaranteed delivery of novel multimedia services over the Internet. Concretely, the thesis consists of five main contributions (also annotated in the figure):

- 1) A framework for negotiating, configuring and maintaining end-to-end federations of network domains for the delivery of Internet-based multimedia services with stringent QoS requirements.
- 2) A novel cache replacement strategy for multimedia content. Traditional cache replacement strategies directly apply historical information to decide what content to cache. In contrast, the presented strategy uses predicted information about future popularity instead.
- 3) A hierarchical network management architecture to support scalable communication and collaboration between AEs, both within and across network domains.
- 4) A substrate to facilitate dissemination of context information between AEs, called the Semantic Communications Bus (SCB). It is augmented with semantics, allowing the SCB to intelligently filter relevant information based on meaning, rather than syntactic characteristics.
- 5) An algorithm for semantic matchmaking and discovery of management services. In addition to exchanging information, AEs interact by instantiating specialized management functionality offered by other AEs. The algorithm matches semantic descriptions of management functionality with AE goals and requirements.

The first two contributions are related to the federated network management part of our work and are described in more detail in Section III. The latter three are related to the autonomic network management component and are discussed in Section IV.

A. Publications

A digital version of the thesis book is available on-line [5]. The results of this research have been published several international journals and conference proceedings. Specifically, the research lead to 5 journal publications [6], [7], [8], [9], [10]. These articles were published in Elsevier Computer Communications [6], the Journal of Communications and Networks [7], the Wiley International Journal of Network Management [8], [9] and the Elsevier Journal of Network and Computer Applications [10]. Additionally, the research lead to the publication of 13 papers in the proceedings of international network and service management conferences [11], [12], [13], [14], [15], [16], [17], [18], [19] and workshops [20], [21], [22], [23], such as MANWEK, IM, NOMS, ManFI, MACE and DANMS.

B. Collaborations

The presented research was conducted in collaboration with partners from both academia and industry. Specifically, parts of this research were carried out in the context of 2 European projects: CELTIC RUBENS and FP7 STREP OCEAN. Additionally, several publications are the result of collaboration with prof. John Strassner.

III. FEDERATED NETWORK MANAGEMENT

A. Federated resource management

End-to-end QoS guarantees over the Internet can only be achieved through cooperation of the network domains on the end-to-end route through the Internet between the multimedia content provider and its consumers. The dissertation presents a framework to automatically negotiate federation agreements, which supports such cooperation [19], [23]. Specifically, the framework sets up federations between the multimedia content provider and a set of intermediary core Internet domains. This facilitates the reservation of QoS-guaranteed end-to-end paths towards the consumers. The consumers are access ISPs, which deliver the content over a managed IP network to their own customers, a set of end-users. The federation can additionally incorporate cloud storage providers. This allows the content provider to dynamically reserve storage resources, and deploy content caches. These caches reduce the total bandwidth consumption and thus decrease delivery costs. Additionally, as these caches are deployed inside the network, they can be shared among several consumers. Cache sharing increases the caches efficiency, without increasing the total storage cost.

Specifically, the process of setting up multimedia delivery federations is achieved through a series of interactions between the involved stakeholders (i.e., access ISP, multimedia content provider, core internet network domain and cloud storage provider). The details of the interactions between the access ISP and content provider are depicted in Figure 2. As shown, the access ISP initiates the federation formation procedure by contacting the multimedia content provider (CP). In return the CP replies with a set of available QoS levels for the delivery of its content. Subsequently, the ISP requests an initial price offer for one or more QoS levels, which is followed by an optional price negotiation process. Finally, if the negotiation ends in agreement, the CP includes the necessary partners (i.e., cloud




Fig. 2. A sequence diagram detailing the negotiation process between the content provider and an access ISP customer

storage providers and intermediary ISPs) in the federation and reserves the required amount of resources.

Obviously, determining the price for delivering content with specific QoS guarantees and determining which network and storage resources should be allocated to meet them, is far from trivial. To accomplish this, we present an algorithm that finds the most cost effective set of federation partners and accompanying resources offered by these partners, in order to satisfy requested QoS. Specifically, the algorithm consists of 4 steps and iteratively improves an initial solution. Figure 3 presents a flowchart of the algorithm's 4 steps. In the first step, the algorithm calculates the cheapest direct path (i.e., containing only intermediary ISPs and no cloud storage providers) that satisfies the end-to-end QoS requirements from the content provider to the access ISP. This problem can be directly translated into the NP-complete multi-constrained optimal path (MCOP) problem [24], [25]. As such, an existing MCOP algorithm or heuristic can be employed to solve this sub-problem. Subsequently, the algorithm can iteratively improve this initial solution. This improvement is done by adding cloud storage providers to the path. During each iteration, one additional cloud storage provider is added to the path. Using linear optimization techniques, the optimal cost-minimizing storage capacity that should be reserved with each provider is calculated. Throughout the third step, the algorithm identifies opportunities to combine the delivery paths of several access ISPs. If they contain the same cloud storage site, the deployed caches can be shared, which could potentially reduce the storage costs associated with both ISPs. Finally, in the fourth step, the algorithm finds the cheapest path for each access ISP, from all the different alternative solutions that have been calculated in steps 2 and 3.

The proposed algorithm was thoroughly evaluated, based on simulation results. The evaluation showed that including cloud storage providers into the federations significantly reduces the delivery costs compared to traditional end-to-end QoS reservation mechanisms that use only direct QoS-constrained paths between the content provider and its customers. Additionally, cache sharing was proven to reduce the delivery costs even further for access ISPs positioned near each




Fig. 3. Flowchart depicting the steps and flow of the resource reservation algorithm

other in the network.

B. Predictive multimedia content caching

The federated service delivery framework, presented in Section III-A, employs dynamically deployed caches to reduce the considerable bandwidth requirements and associated costs of modern multimedia services. A cache transparently stores a subset of the content closer to the end-users, allowing it to be served without having to traverse the entire network and thus reducing bandwidth requirements and transmission costs. As caches can only store a subset of the available content, a cache replacement strategy is needed to determine what content to cache. Traditional strategies directly use the historical popularity as an indicator of future popularity. In contrast, we propose a novel strategy that first predicts the future popularity of content, taking into account popularity dynamics [17], [10]. Additionally, an algorithm is presented to perform the actual popularity predictions.

The proposed content popularity prediction algorithm estimates the evolution of a multimedia content item's cumulative request pattern, based on historical information. It achieves this by using the Levenberg-Marquardt [26] non-linear optimization technique to fit a set of models to the available historical input data. These fitted models can subsequently be used to extrapolate the request pattern's future evolution. Finally, a selection metric is used to select the optimal model that will be used as a predictor. Based on the request patterns of a real Video on Demand (VoD) service, we selected four popularity models to be used by the algorithm: constant, power-law, exponential and gaussian. The algorithm is graphically illustrated by way of an example in Figure 4. It shows the request trace of an actual video in a deployed VoD system over the course of 448 hours. The prediction algorithm was applied to the first 400 hours of the trace (the known history) using the exponential and gaussian distributions. The parts of the curves before the vertical line represent the fits to known history, while the parts after the line represent the predictions. The actual number of requests that occur in the interval $[400, 448[$ is 15. The exponential distribution predicts 14.55 requests within that interval, which is very close to the real value. On the other hand, the gaussian distribution predicts only 1.25. This is also reflected in the figure, which shows that the gaussian distribution poorly approximates the start and end of the request pattern.




Fig. 4. Optimal fit of the exponential and gaussian distributions to an example cumulative request pattern; the vertical line represents the current point in time t , on its left is the 400-hour known history, on its right the 48-hour predicted future




Fig. 5. Comparison of the different cache replacement strategies in terms of cache hit rate, as a function of cache size

In our work, we propose a cache replacement strategy based on this popularity prediction strategy. It is called Predictive Least Frequently Used (P-LFU), and is a predictive version of the traditional LFU caching strategy. LFU keeps track of the number of times that every object is requested. The objects that receive the highest number of requests within a specified time frame (i.e., the history window) are kept in cache [27]. In contrast, the P-LFU strategy uses the predicted number of requests for each object, instead of the known number of requests in the past.

Both the proposed popularity prediction algorithm and the predictive cache replacement strategy were evaluated on the previously mentioned real VoD request traces. The novel P-LFU strategy was compared to traditional cache replacement strategies, such as LFU and Least Recently Used (LRU), as well as the theoretical optimal MIN strategy [28], [29]. The results, as shown in Figure 5, prove that, in terms of cache hit rate, P-LFU in combination with perfect predictions (i.e., PP-LFU) performs up to 20% better than LFU in the evaluated scenario. Additionally, P-LFU in combination with the presented popularity prediction algorithm (i.e., OP-LFU) has a gain of up to 10%.

IV. AUTONOMIC NETWORK MANAGEMENT

The previous section described our proposed federated service delivery framework, which is capable of negotiating and configuring Internet-wide network federations. Once such




Fig. 6. An overview of the collaborative, AE-driven, hierarchical management architecture

a dynamic federation of large-scale network domains has been configured, it needs to be effectively managed. To achieve this in a scalable manner, the self-governing capabilities of the network should be distributed across many autonomic management components, or AEs. These AEs should be able to efficiently interact, communicate and collaborate. The contributions of the dissertation presented in this section aim to achieve this.

A. Hierarchical management architecture

Traditional network management frameworks are centralized and monolithic entities, which makes them ill-suited to cope with the ever-increasing growth of modern communications networks. It has been generally agreed upon that an autonomic management framework for managing the Future Internet should be highly distributed to maintain scalability [30], [31], [32]. Nevertheless, the flat distributed approach, where all components act as equals, has several disadvantages. Specifically, it maps poorly to the hierarchical management responsibilities and limits the view of AEs on the managed environment resulting in sub-optimal decision making. Recently, an alternative distributed approach has been suggested, where AEs are structured in a hierarchy [31], [33]. In our work, we build on this idea and propose a hierarchical management structure for autonomic network management components [21], [7]. This allows the network overhead associated with managing resources to be greatly reduced. Additionally, the dissemination of context, propagation of policies and collaboration between AEs can be more efficiently orchestrated, resulting in a more scalable management architecture.

Figure 6 gives an overview of how the AEs interact within and across network domains. AEs at the bottom of the hierarchy interact directly with the underlying infrastructure. They gather monitoring information and are responsible for directly configuring managed resources. The bottom level AEs process and annotate the gathered information and convert it to semantically annotated context. Subsequently, the semantic information is summarized and aggregated, after which it is propagated upwards through the hierarchy. Additionally, the parent AE can install semantic filters, which allow it to receive only the information that it needs. These filters can be dynamically and automatically adapted, as context requirements may change based on the state of the managed environment [34]. The goal of the aggregation and filtering process is to maintain scalability throughout the hierarchy. It




Fig. 7. The amount of context received by datacenter level AEs, for multiple datacenter hierarchy levels

gives AEs at higher levels a broader, but less detailed view on the managed environment, while giving low level AEs a very precise but narrow view. This allows the distributed management framework to perform high-level decisions and optimizations, as well make detailed configuration changes. More details about the semantic context filtering and aggregation process are presented in Section IV-B.

In line with the upward stream of context information, there is a downward stream of management decisions and policies. The high-level management decisions taken at the top of the hierarchy, are converted into more concrete form as they propagate downwards. At the bottom, only concrete device configuration changes remain. Similarly, management policies need to be adapted to be in line with the knowledge and responsibilities of the AEs to which they are propagated. At the top of the hierarchy, policies represent the high-level business goals of the network domain's operator, while at the bottom they represent specific resource policies. This maps well to the Policy Continuum [35].

Finally, the management of federations is supported by way of peer relationships among different management hierarchies. High-level AEs of different domains communicate in order to negotiate contracts and configure federation agreements. To be able to set up and maintain such inter-domain collaborations, these AEs will also need to exchange context information. In contrast to context exchange within a domain, inter-domain context exchange is complicated by security and trust issues among operators.

The quantitative merits of the proposed hierarchical approach were evaluated and compared to the flat distributed approach using an analytical formulation. Figure 7 depicts analytical results for an AE hierarchy managing a cloud datacenter as a function of the number of servers within the datacenter. The hierarchy consists of 1 level of server AEs at the bottom, 2 or 3 levels of datacenter AEs (each managing a cluster of server AEs or a cluster of lower level datacenter AEs) and a single root cloud AE at the top. The results in Figure 7a show the amount of context that is received per time unit by each datacenter AE for 2 levels of datacenter AEs, while Figure 7b presents the same results for 3 levels. The figures show that the top datacenter level shows a linear increase in the amount of receive context information as a function of the number of servers. This obviously presents a scalability issue. However, as shown in Figure 7b, the slope of this linear increase is much lower if an additional level is added to the hierarchy. Additionally, the context that needs to be processed




Fig. 8. The SCB plays a central role in the interactions between AEs and network resources; its core ontologies are used in the matchmaking and context dissemination processes

by lower levels is independent of the number of servers. As such, the study allows us to conclude that scalability of the management framework can be maintained by increasing the number of levels within the hierarchy together with the amount of managed resources.

B. Semantic context dissemination

As stated, the hierarchical management framework copes with the ever-increasing amount of management information by intelligently filtering and aggregating it. In the dissertation, the Semantic Communications Bus (SCB) is proposed to perform this task [22], [8]. AEs specify their interests in information via semantic filter rules. The SCB employs semantic reasoning to match information with these rules. The use of semantics guarantees the unambiguous interpretation of exchanged information. This is needed when AEs need to take management decisions based on the information they receive. Additionally, semantic interpretation of information is an important aspect of communication among federation partners [36]. The SCB achieves semantic consistency among AEs through a set of shared ontologies. These ontologies semantically model the managed environment and semantically represent context information and filter rules.

Figure 8 shows how the SCB is used for the dissemination of context. It additionally depicts the semantic matchmaking process, which will be explained in Section IV-C. Specifically, it uses the publish/subscribe paradigm. An AEs context manager may decide that in order to perform its management functions, it requires specific types of context information. It constructs a filter rule, which represents this type of information, and sends it to the SCB. When the SCB receives a chunk of context from another AE, it uses semantic reasoning to determine if the information matches the filter rule. If it does, the information is forwarded to the requesting AE.

The SCB supports three types of filter rules, in the form of OWL class expressions [37], SWRL rules [38], or Jena rules [39]. Context information is expressed as OWL class instances. Matching with OWL filter rules is as simple as asking the OWL reasoner if the message instance belongs to the filter rule OWL class. For SWRL and Jena rules the process is slightly more complex. There, the rule consequent asserts that the context message belongs to the subscriber that inserted the rule. As such, if a context instance satisfies the antecedent of a

rule, it is asserted that it belongs to the associated subscriber. Subsequently, the OWL reasoner can be used to retrieve all context that belongs to a subscriber.

As part of the dissertation, a prototype of the SCB was implemented and integrated into the hierarchical AE architecture. An evaluation of the integrated solution was performed, using the management of large-scale federated cloud datacenters as a use-case [7]. This showed the viability of using semantic reasoning to intelligently disseminate context in the face of a huge amount of generated monitoring information.

C. Semantic service matchmaking

Through the dissemination of context, AEs can detect problems in the networks state. However, they may not always be able to solve the detected problems themselves. Therefore, to further support loosely-coupled collaboration between AEs, a matchmaking mechanism for the dynamic discovery of management services and functions is needed. Matchmaking is defined as the process of discovering a set of services or functions that fulfil a given set of requirements. The proposed matchmaking algorithm is capable of matching AE goals with management functions offered by other AEs.

Figure 8 depicts how AEs use the semantic matchmaking functions of the SCB. AEs can publish their semantic service descriptions. Similarly to the context dissemination process, these descriptions use the concepts defined in the shared ontologies and are defined using SWRL atoms. A description contains inputs, outputs, preconditions and effects (IOPEs). The inputs represent the context that is required by the management service in order to perform its task. On the other hand, the outputs depict context that will be generated by the management service. The preconditions define the environmental conditions that must hold for the service to be usable, while the effects describe the conditions that will hold after the services execution. Together, the preconditions and effects thus describe how the managed environment will change when the management service is executed. This allows AEs to better match these services with their planned actions. Requested functionality is described in a similar fashion.

For a described and requested function to be compatible, they do not need to be semantically identical. Instead, the IOPEs need to satisfy a set of subsumption relationships. An ontological concept D subsumes a concept B if B is a specialization of D (or alternatively, if D is a generalization of B). As such, the dissertation presents a matchmaking algorithm that checks the subsumption relationships between the IOPEs of descriptions. It is capable of determining semantic subsumption between any kind of SWRL atoms. An evaluation of the SCB prototype implementation showed that the algorithm is capable of matching over 200 service descriptions in less than 1 second on commodity hardware.

V. CONCLUSION

In our work, we proposed a federated architecture to autonomously manage multimedia delivery networks in a scalable manner. The first main contribution consists of a framework to negotiate and configure federations between the stakeholders involved in the delivery of multimedia services. An algorithm is presented that identifies the optimal set of

federation partners and resources that should be shared among them. The second main contribution consists of an architecture that structures autonomic management components in a hierarchy. It allows them to exchange context information in a scalable manner through the use of semantic filtering and aggregation. Additionally, it supports the discovery of management functionality through the semantic description of management goals.

Through simulation, analytical studies and prototype implementations, the proposed solutions were thoroughly studied. These studies showed that including additional actors, such as cloud storage providers, in a multimedia delivery federation can greatly increase efficiency and decrease costs. Additionally, the hierarchical management architecture is proven to greatly increase scalability and management optimality compared to centralized or flat distributed architectures. Finally, the prototype showed the viability of using semantic context dissemination and service matchmaking in a large-scale management scenario.

ACKNOWLEDGMENT

Jeroen Famaey's PhD was funded by the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT-Vlaanderen) under grant no. 73185.

REFERENCES

- [1] N. Agoulmine, Ed., *Autonomic network management principles: From concepts to applications*. Elsevier, 2011.
- [2] L. G. Roberts, "A radical new router," *IEEE Spectrum*, vol. 46, no. 7, pp. 34–39, 2009.
- [3] B. Jennings, K. C. Feeney, R. Brennan, S. Balasubramaniam, D. Botvich, and S. van der Meer, "Federating autonomic network management systems for flexible control of end-to-end communications services," in *Autonomic network management principles: From concepts to applications*, N. Agoulmine, Ed. Elsevier, 2011, pp. 101–118.
- [4] B. Jennings, S. van der Meer, S. Balasubramaniam, D. Botvich, M. Ó Foghlú, W. Donnelly, and J. Strassner, "Towards autonomic management of communications networks," *IEEE Communications Magazine*, vol. 45, no. 10, pp. 112–121, 2007.
- [5] J. Famaey, *Federated and Autonomic Management of Multimedia Services*, 2012, available online: http://users.ugent.be/~jfamaey/phd_jfamaey.pdf, ISBN: 978-90-8578-523-1.
- [6] J. Famaey, T. Wauters, F. De Turck, B. Dhoedt, and P. Demeester, "Network-aware service placement and selection algorithms on large-scale overlay networks," *Computer Communications*, vol. 34, no. 15, pp. 1777–1787, 2011.
- [7] J. Famaey, S. Latré, J. Strassner, and F. De Turck, "A hierarchical context dissemination framework for managing federated clouds," *Journal of Communications and Networks*, vol. 13, no. 6, pp. 567–582, 2011.
- [8] —, "Semantic context dissemination and service matchmaking in future network management," *International Journal of Network Management*, vol. 22, no. 4, pp. 285–310, 2012.
- [9] J. Famaey and F. De Turck, "Federated management of the future internet: Status and challenges," *International Journal of Network Management*, vol. 22, no. 6, pp. 508–528, 2012.
- [10] J. Famaey, F. Iterbeke, T. Wauters, and F. De Turck, "Towards a predictive cache replacement strategy for multimedia content," *Journal of Network and Computer Applications*, 2012, accepted.
- [11] J. Famaey, T. Wauters, F. De Turck, B. Dhoedt, and P. Demeester, "Towards efficient service placement and server selection for large-scale deployments," in *4th Advanced International Conference on Telecommunications (AICT)*, 2008, pp. 13–18.

- [12] —, “Dynamic overlay node activation algorithms for large-scale service deployments,” in *19th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM)*, 2008, pp. 14–27.
- [13] J. Famaey, W. De Cock, T. Wauters, F. De Turck, B. Dhoedt, and P. Demeester, “A latency-aware algorithm for dynamic service placement in large-scale overlays,” in *11th IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2009.
- [14] J. Famaey, J. Donders, T. Wauters, F. Iterbeke, N. Sluijs, B. De Vleeschauwer, F. De Turck, P. Demeester, and R. Stoop, “Comparative study of peer-to-peer architectures for scalable resource discovery,” in *First International Conference on Advances in P2P Systems (AP2PS)*, 2009.
- [15] J. Famaey, W. Van de Meerssche, S. Latré, S. Melis, T. Wauters, F. De Turck, K. De Schepper, B. De Vleeschauwer, and R. Huysegems, “Towards intelligent scheduling of multimedia content in future access networks,” in *12th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2010.
- [16] H. Moens, J. Famaey, S. Latré, B. Dhoedt, and F. De Turck, “Design and evaluation of a hierarchical application placement algorithm in large scale clouds,” in *12th IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2011.
- [17] J. Famaey, T. Wauters, and F. De Turck, “On the merits of popularity prediction in multimedia content caching,” in *12th IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2011.
- [18] K. De Schepper, B. De Vleeschauwer, C. Hawinkel, W. Van Leekwijck, J. Famaey, W. Van de Meerssche, and F. De Turck, “Shared content addressing protocol (scap): Optimizing multimedia content distribution at the transport layer,” in *13th IFIP/IEEE Network Operations and Management Symposium (NOMS)*, 2012.
- [19] J. Famaey, S. Latré, T. Wauters, and F. De Turck, “Fedrr: A federated resource reservation algorithm for multimedia services,” in *13th IFIP/IEEE Network Operations and Management Symposium (NOMS)*, 2012.
- [20] J. Famaey, B. De Vleeschauwer, T. Wauters, F. De Turck, B. Dhoedt, and P. Demeester, “Dynamic QoE optimisation for streaming content in large-scale future networks,” in *1st IFIP/IEEE International Workshop on Management of the Future Internet (ManFI)*, 2009.
- [21] J. Famaey, S. Latré, J. Strassner, and F. De Turck, “A hierarchical approach to autonomic network management,” in *2nd IFIP/IEEE International Workshop on Management of the Future Internet (ManFI)*, 2010.
- [22] —, “An ontology-driven semantic bus for autonomic communication elements,” in *5th International Workshop on Modelling Autonomic Communication Environments (MACE)*, 2010.
- [23] J. Famaey, S. Latré, T. Wauters, and F. De Turck, “An sla-driven framework for dynamic multimedia content delivery federations,” in *5th International Workshop on Distributed Autonomous Network Management Systems (DANMS)*, 2012.
- [24] T. Korkmaz and M. Krunz, “Multi-constrained optimal path selection,” in *Proceedings of the Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2001, pp. 834–843.
- [25] H. Pouyllau and R. Douville, “End-to-end QoS negotiation in network federations,” in *Proceedings of the 12th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2010, pp. 173–176.
- [26] K. Levenberg, “A method for the solution of certain non-linear problems in least squares,” *Quarterly Journal of Applied Mathematics*, vol. 2, no. 2, pp. 164–168, 1944.
- [27] N. Megiddo and D. S. Modha, “Outperforming LRU with an adaptive replacement cache algorithm,” *Computer*, vol. 37, no. 4, pp. 58–65, 2004.
- [28] J. Gecsei, D. R. Slutz, and I. L. Traiger, “Evaluation techniques for storage hierarchies,” *IBM Systems Journal*, vol. 9, no. 2, pp. 78–117, 1970.
- [29] B. Van Roy, “A short proof of optimality for the MIN cache replacement algorithm,” *Information Processing Letters*, vol. 102, no. 2–3, pp. 72–73, 2007.
- [30] Y. Cheng, R. Farha, M. S. Kim, A. L. Garcia, and J. W. K. Hong, “A generic architecture for autonomic service and network management,” *Computer Communications*, vol. 29, no. 18, pp. 3691–3709, 2006.
- [31] J. Strassner, J. Won-Ki Hong, and S. van der Meer, “The design of an autonomic element for managing emerging networks and services,” in *IEEE International Conference on Ultra Modern Telecommunications (ICUMT)*, 2009.
- [32] L. Baresi, A. D. Ferdinando, A. Manzalini, and F. Zambonelli, “The CASCADAS framework for autonomic communications,” in *Autonomic Communication*, 2009, pp. 147–168.
- [33] J. Strassner, N. Agoulmine, and E. Lehtihet, “FOCALE – a novel autonomic networking architecture,” in *Latin American Autonomic Computing Symposium (LAACS)*, 2006.
- [34] S. Latré, S. van der Meer, F. De Turck, J. Strassner, and J. Won-Ki Hong, “Ontological generation of filter rules for context exchange in autonomic multimedia networks,” in *12th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2010, pp. 575–582.
- [35] J. Strassner, *Policy-Based Network Management – Solutions for the Next Generation*. Morgan Kaufman, 2004.
- [36] K. Feeney, R. Brennan, J. Keeney, H. Thomas, D. Lewis, A. Boran, and D. O’Sullivan, “Enabling decentralised management through federation,” *Computer Networks*, vol. 54, no. 16, pp. 2825–2839, 2010.
- [37] W3C, “OWL 2 web ontology language document overview,” Online, 2009, <http://www.w3.org/TR/owl2-overview>.
- [38] I. Horrocks, P. F. Patel-Schneider, H. Boley, S. Tabet, B. Grosz, and M. Dean, “SWRL: A semantic web rule language combining OWL and RuleML,” Online, 2004, <http://www.w3.org/Submission/SWRL/>.
- [39] J. J. Carroll, I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne, and K. Wilkinson, “Jena: Implementing the semantic web recommendations,” in *13th International World Wide Web Conference*, 2004.