Self-management of Hybrid Networks: Introduction, Pros and Cons

Tiago Fioreze and Aiko Pras

University of Twente
Enschede, The Netherlands
t.fioreze@utwente.nl, a.pras@utwente.nl

Abstract. In the last decade 'self-management' has become a popular research theme within the networking community. While reading papers, one could get the impression that self-management is the obvious solution to solve many of the current network management problems. There are hardly any publications, however, that discuss the drawbacks of self-management. In this paper we will therefore introduce self-management for the specific case of hybrid networks, and discuss some pros and cons. In particular, this paper investigates the feasibility to employ self-management functions within hybrid optical and packet switching networks. In such networks, large IP flows can be moved from the IP level to the optical level, in an attempt to reduce the load at the IP layer and enhance the quality of service (QoS) of the flow that is moved to the optical level. One of the typical management tasks within such networks, is the establishment and release of lightpaths. This paper identifies the advantages and disadvantages of introducing self-management to control such lightpaths.

Keywords: Future Internet, self-management, hybrid networks

1 Introduction

In recent years there has been a considerable interest in what is called the *Future Internet*. Two fundamentally different approaches are under discussion: evolution versus revolution. The evolutionary approach aims at moving the Internet from one state to another through incremental patches. The revolutionary approach, on the other hand, proposes a radical redesign of the current Internet architecture, and is therefore also called *clean-slate* approach [2].

Whatever approach will prevail, we can already foresee a future Internet in which optical communication will play a major role. At this moment we can already observe that the core Internet, which once solely relied on IP routing to deliver end-to-end communications, is moving towards a hybrid optical-IP network. Such network takes data forwarding decisions simultaneously at both IP and optical level [11]. It is composed of intermediate multi-service devices that are both switches at the optical level and traditional routers at the IP level.

In such an environment, data flows can traverse a hybrid network through either an IP path or a lightpath. In this paper, we consider a lightpath as a direct connection over an optical fiber; the lightpath can consist of the whole fiber, a wavelength within the fiber (lambda), or a TDM-based channel within a lambda.

Traditionally, the establishment and release of lightpaths is controlled by human managers. In this article we investigate the feasibility of removing the human manager from the loop and to introduce self-management capabilities. Such capabilities enable an autonomic decision process to configure lightpaths, based on measurement data received from the hybrid network. The human manager expresses *what* the self-managing system is expected to achieve, but not *how* this should be done. The human manager is therefore moved to a higher level in the management hierarchy, where he controls the autonomic decision process, rather than the whole hybrid network. Figure 1 shows this self-management approach.

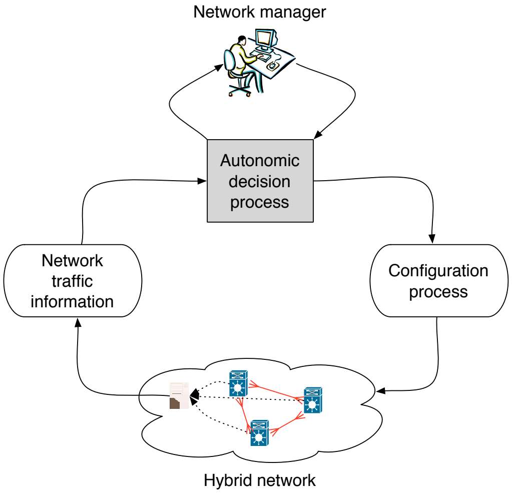

Fig. 1. The self-management approach.

The goal of this article is to introduce, in a tutorial style, the main advantages and disadvantages of self-management within hybrid networks. In fact, the material presented in this article summarizes four years of PhD research; readers interested in the technical details of our approach and the validation behind the conclusions, are encouraged to read the thesis and associated papers [3] [6] [7] [8].

The remainder of this article is structured as follows. Section 2 introduces the concept of self-management, and provides some definitions. Following that, Section 3 identifies the related work in this area. Next, Section 4 introduces our approach to employ self-management in hybrid networks. Sections 5 and 6 then respectively present the pros and cons of employing our self-management proposal in hybrid networks. Finally, in Section 7 we draw our conclusions.

2 What is self-management?

Self-management encompasses the act of computer systems managing their own operation without (or with very little) human intervention. It was first defined by IBM in 2001 with the IBM Autonomic Computing Initiative (ACI) manifesto [9]. In such manifesto, IBM proposed an approach in which self-managed computing systems could work with a minimum of human interference. This approach is inspired from the human body's autonomic nervous system. Many actions are performed by our nervous system without any conscious recognition, such as the act of adjusting our eye's pupils depending on the amount of light or the act of sweating in order to regulate our body temperature. Below, we quote the main objective of IBM's autonomic initiate that is:

> "*to design and build computing systems capable of running themselves, adjusting to varying circumstances, and preparing their resources to handle most efficiently the workloads we put upon them. These autonomic systems must anticipate needs and allow users to concentrate on what they want to accomplish rather than figuring how to rig the computing systems to get them there.*"

A system can be seen as a collection of computing resources bound together to achieve certain objectives. For example, a network router can constitute a system responsible for forwarding network traffic. When combined with other network routers, they can form a larger system, *i.e.*, a Local Area Network (LAN). On its turn, a LAN network combined with other LANs can form a Metropolitan Area Network (MAN), and so on. Based on the IBM autonomic principle, each system must be able to manage its own actions (*e.g.*, traffic forwarding), while collaborating with a larger, higher-level system. The same analogy can be found in the human body. From single cells to organs and organ systems (*e.g.*, the circulatory system), each level maintains a measure of independence while contributing to a higher level of organization, culminating in the organism, *i.e.*, the human body. In most parts of our daily life, we remain unaware of our vital organs (*e.g.*, the heart) activities, since these organs (systems) take care of themselves and they only ascend to a higher level (*e.g.*, the brain) when something is wrong and they need some assistance.

2.1 Self-management aspects

IBM divided self-management into 4 aspects (yet other subdivisions exist), commonly referred as *self-**, as follows:

- *Self-configuration*: consists of an automated configuration process of components and systems based on high-levels policies. For example, when a new device is taken into a network, this device is expected to automatically configure itself and at the same time the rest of the network seamlessly adjust itself to incorporate this new device.

- *Self-optimization*: means that components and systems are supposed to continuously improve their own performance. One example of this aspect is the automatic update process most operating systems provide to their users. Instead of requiring users to manually seek for updates, the operating system does that automatically.
- *Self-healing*: consists of the capability of a system to automatically detect, diagnose, and repair problems found at certain components. As an example, a computer could self-heal every time a virus would strike the system, by automatically patching the damaged files.
- *Self-protection*: is seen as a system automatically defending itself against malicious attacks or failures. A computer system could, for instance, prevent the infection by a certain email virus through analysis of email attachments.

2.2 Different definitions for self-management

Although the term self-management has been widely used, there is no universal consensus among authors on what self-management actually means, which leads to different definitions for the term self-management. Some of the most known definitions for self-management are as follows:

- *Autonomic management*: is the most common synonym used to refer to the term self-management. That comes from the fact IBM considers self-management as the essence for autonomic computing systems. As a result, the terms self-management and autonomic management are interchangeably used to mean the same. By analyzing the keywords attached to papers submitted to important network management conferences (*e.g.*, IM, NOMS, CNSM), we found out that 80% of the papers were submitted with the keywords as self-*, whereas 20% were registered as autonomic. This leads us to a conclusion that even if they were constantly used as synonyms, the term self-management seems to be the most referred and used by the network management community.
- *Automatic management*: is commonly confused as autonomic management (and thus with self-management). Automatic management refers as the act of managed devices automatically following explicit policies defined by a network operator. In its turn, autonomic management refers as a specialized automatic process in the sense that the process is instructed to perform actions based on certain policies too, **but** with the capability of self-learning new actions.
- *Autonomous management*: is another definition referring to self-management. Autonomous means that a process can operate independently from any human intervention. However, this lack of external control is, according to some, a contradiction. If an autonomous “management” system includes enough intelligence in order for the system to govern its own management, one can assume that there is no need whatsoever of managing such a system, which somehow invalidates the use of the term management to address this kind of management approach.

It is worth saying that the foregoing differentiation among self-management definitions is not a common view in the community. On the contrary, this differentiation solely destined for being a reference to be used throughout this article. We see these definitions as following an evolution in the network management approaches as well as having different degrees of autonomy (Figure 2).

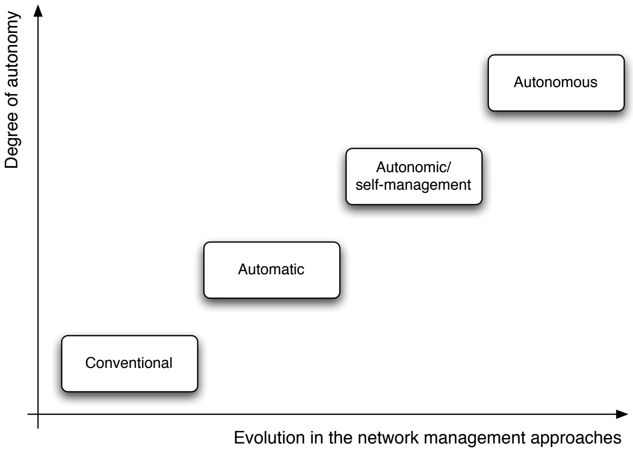

Fig. 2. Evolution in the network management approaches *vs.* their degree of autonomy.

The simplest management approach is the conventional management approach. In the conventional management approach, the network management system is manually managed by network operators. There is no intelligence whatsoever and no (or very little) automation in the execution of management tasks. A next step in the evolution of management approaches is the automation of management tasks. In this case, the management system automatically performs explicit tasks defined by network managers, but nothing beyond the scope of the defined rules. Following to automatic management, autonomic management (or self-management) also performs these tasks, but it is capable of learning new rules by itself. The last step in the evolution process and the most complex one is the autonomous management. At this level, the management system is fully capable of deciding by itself the rules to follow. There is therefore no dependence on human intervention. The management system is intelligent enough to decide its own rules and following them according to its judgement.

3 Related work on self-management

Since the releasing of the ACI manifesto, several research works about the use of self-management have been reported. To name a few of these works, Lupu *et al.* [12] have been researching the use of self-management on healthcare practicing, in which a ubiquitous self-managed computing environment is used to monitor and report the health of patients under medical treatment. In another work, self-management is investigated to be used in situations where there is a great risk for human beings, such as in military or disaster scenarios. Within this line of research, we point out the work by Asmare *et al.* [1] who has been investigating the use of self-management on Unmanned Autonomous Vehicles (UAVs).

Not much differently, self-management has also being investigated in the area of communication networks [10]. Much of the focus of this investigation aims at developing highly distributed algorithms, with the objective to optimize several aspects of network operability (e.g., performance). This optimization is aimed through the provision of self-management capabilities to communication networks. Studies on self-management is also the focus of several research projects, such as Autonomic Internet (AUTOI), Self-Optimisation and self-ConfiguRATion in wirelEss networkS (SOCRATES), and UniverSelf.

A study that is closely related to ours is by Miyazawa et al. [13]. In their research, they propose a dynamic bandwidth control management mechanism based on the volume of IP flows. In their work there is a centralized management system that observes the bandwidth of IP flows, and decides about offloading these flows based on pre-defined upper and lower threshold values. These threshold values are defined in advance by a human operator and statically stored in the configuration file of the management system. Once an IP flow has a bandwidth utilization that exceeds the pre-determined upper threshold, the management system triggers an action to create a lightpath. In contrast, when the flow decreases its bandwidth utilization below the lower threshold, the management system initiates a deletion process for deleting the established lightpath.

4 Self-management of lightpaths

We focus now the use of self-management in the context of hybrid optical and packet switching networks, more specifically on the self-management of lightpaths in these networks. The use of self-management is aimed here at autonomically: 1) detect flows at the IP level eligible to be moved to the optical level as well as 2) establish/release lightpaths for those flows. In this paper we adopt the definition of flows as described in the information model for the IP Flow Information eXport (IPFIX) protocol (RFC 5102). In this RFC, an IP flow is defined as a unidirectional sequence of packets that share the same properties (e.g., the same source and destination IP addresses, source and destination port numbers and higher level protocol).

Network operators are only required to initially configure the self-management process with decision policies. After this initial setup, the self-management process autonomically runs by itself. Decision policies define a desired objective, which must be achieved by the self-management functions. In our research, the main objective is to offload as much traffic as possible from the IP level to the optical level. For that, our self-management approach aims at moving flows to the optical level that are few in amount, but represent most of the traffic, namely the elephant flows. Figure 3 depicts our approach for the self-management of lightpaths in hybrid networks.

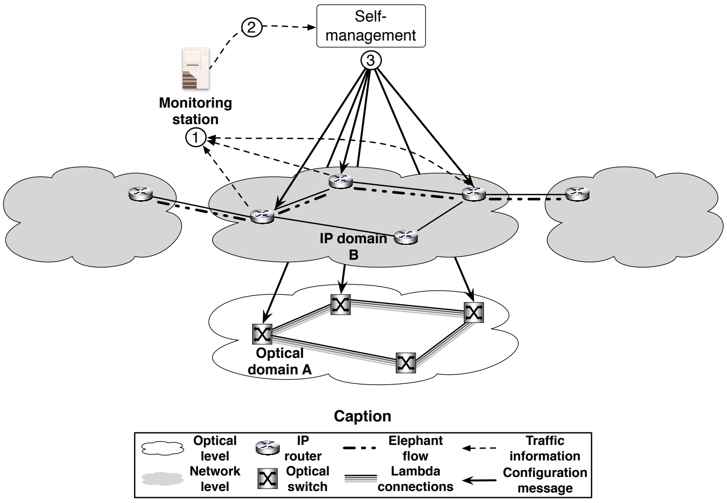

Fig. 3. Self-management of lightpaths in hybrid optical and packet networks.

In Figure 3, IP routers located in the IP domain B are exporting network traffic information to a monitoring station (step 1). Network information contains flow information, such as source & destination IP addresses, protocol, flow volume, amongst others. This information is then forwarded to our self-management module (step 2). Based on the information received, decisions are made by the module taking into account whether an elephant flow is eligible or no longer eligible for a lightpath at the optical level. If the decision is in favor of creating a lightpath (i.e., the elephant flow is eligible to be moved to the optical level), the self-management module configures the IP routers in the IP domain B and the optical switches in the optical domain A (step 3). The routers are informed that the elephant flow is offloaded to the optical level. On their turn, the optical switches are configured to establish a lightpath for the offloaded elephant flow. From that point on, the elephant flow is switched at the optical level bypassing thus the network level in the IP domain B. For configuring routers and switches, existing management technologies can be used, such as the Command Line Interface (CLI), the Generalized Multi-Protocol Label Switching (GMPLS) protocol, the Simple Network Management Protocol (SNMP) or the emerging Network Configuration (NetConf) protocol.

Note that this article can only summarize the operation of our self-management approach; details of this operation can be found in other papers [4] [5], as well as the thesis that resulted from this research [3].

5 Advantages of self-management

Advantages and disadvantages depend on the context self-management is employed. We highlight in this section the main pros of employing self-management principles in the specific context of hybrid optical and packet switching networks.

5.1 Network performance

The use of self-management improves network performance by automatically reducing the burden of the IP level. When IP flows are completely transported via lightpaths they bypass the per hop routing decisions of the IP level. As a result, the QoS offered by hybrid networks is considerably better when compared to traditional IP networks. Big IP flows that overload the regular IP level, for example, may be moved to the optical level where they experience better QoS (*e.g.*, negligible jitter and larger bandwidth). At the same time, the IP level is offloaded and can better serve smaller flows. Last but not least, it is also cheaper to send traffic at the optical level than at the IP level. For the same traffic rate, the cost of an optical switch is $1/10^{\mathrm{th}}$ of an Ethernet switch or $1/100^{\mathrm{th}}$ of a conventional router.

5.2 Network management

We believe that human factors have an impact on the management of lightpaths. For example, network operators of SURFnet reported, when informally interviewed, that it may take hours (intra-domain) or even days (inter-domain) before a lightpath is established by network operators when using a traditional network management paradigm. In such paradigm, a network manager regularly monitors a hybrid network. Based on his analysis of the collected data, he may decide to establish or release a lightpath. It is worth highlighting that this paradigm keeps the human in the management loop. That is, most of the management decisions have to go through the network manager. As a result, the management system does not go beyond any predetermined state or perform any unexpected action, unless explicitly triggered by the network manager. We argue that during the long periods lightpaths are established, several big IP flows could have been transported via lightpaths, but due to the decision delay they remain being routed at the IP level. Moreover, in such long periods, many large IP flows may be using resources at the IP level and, therefore, likely congesting the IP level. Moreover, by the time the lightpath is finally established, those large flows may no longer exist. The human intervention required to select IP flows and manage lightpaths may be considered therefore slow and inefficient. We see our self-management proposal as an alternative to overcome this dependency on human intervention and therefore improve the network management.

5.3 Selection of unknown large flows

Nowadays, IP traffic from several specialized applications, some of them requiring considerable amounts of bandwidth, already profit from lambda-switched networks capabilities. Examples are: Grid applications, High-Definition Television (HDTV) broadcasting, and large-scale scientific experiments. The knowledge of the heavy-hitter behavior of flows originated from these applications allows network managers to establish lambda-connections in advance for such flows. However, there may be also other big IP flows in current networks that could also benefit from being moved to lambda-connections, but since the network manager is not aware of their existence, they may not be selected. Self-management comes handy here since flows are monitored and selected by the self-managing system rather than by the human manager.

5.4 Dynamic flow selection

The selection of flows to be moved to the optical level are traditionally made based on pre-defined upper and lower threshold values. These threshold values are defined in advance by a human operator and statically stored in the configuration file of the management system (Section 3). The main shortcoming of using thresholds is that they are statically defined and they are not adjusted depending on the current traffic. This can lead up to an unbalance between the IP and optical levels. If the upper threshold values are too restrictive, IP flows may not be offloaded over lightpaths, which may result in congestion in the IP level and underutilization of the optical level. Moreover, with a misadjusted lower threshold, a flow can be inadequately removed from the optical level back to the IP level, where it can contribute to a congestion situation. Our alternative, on the other hand, aims at prioritizing flows by merit (behavior) rather than by characteristics (i.e., port numbers, ToS, and so on). Their merit is measured based on the amount of traffic they are expected to generate. Flows that are expected to generate more traffic are chosen over flows that are expected to generate less traffic.

6 Disadvantages of self-management

Self-management of hybrid networks also introduces a number of problems, as will be explained in this section.

6.1 Complexity in the network management system

As shown in Figure 2, management approaches follow an evolution that is proportional to their degree of autonomy. In the simplest management approach, network managers are responsible for all management tasks. However, as more experience is obtained with these management tasks, some tasks can gradually be automated, which means that the need for human intervention can be reduced. In order to avoid problems related to centralization (single point of failure, possible performance bottleneck), subsequent cycles in the design of the management system may focus on distributing such management tasks. Therefore there is a shift from centralized and explicit management towards distributed and implicit management approaches, such as the aforementioned automatic and autonomic approaches [14]. However, the price to be paid for such evolution is the increased complexity of the management system. The chance that errors get introduced in the implementation of the management system therefore increases, and debugging possible failures becomes harder.

6.2 Network security

Network security consists of providing means to protect network resources from unauthorized access or malicious activities. Consistent and continuos monitoring of network activity is important to prevent or detect any misusage that may fall upon a managed network. Within the context of our self-management approach, network security is expected to prevent any inappropriate use of lightpaths. For instance, Denial of Service (DoS) attacks that may be transiting at the IP level should not be moved to lightpaths. Such move could make the attack more severe, since the increased bandwidth at the optical level allows the transfer of more packets to the attacked system. Other security concepts, such as authorization and intrusion detection should therefore be considered as well, in order to tighten security in hybrid networks. Flows could be verified prior to the offload to the optical level to detect malicious behavior. If flows behave suspiciously, their offload over lightpaths could be blocked and their behavior could be logged for audit purposes and later analysis.

6.3 Temporarily reduction in throughput performance

Another interesting question is whether there is any performance degradation when an active flow is moved from the IP level towards a lightpath. We may expect that, at the moment flows are moved, massive re-ordering takes place, since the first packets transferred over the lightpath can arrive earlier than the last packets over the IP path. To analyze this effect, we used ns-2 to simulate the behavior of TCP flows during such movement, and identified which factors limited the throughput of such flows [15]. For this analysis we used the network topology as shown in (Figure 4).

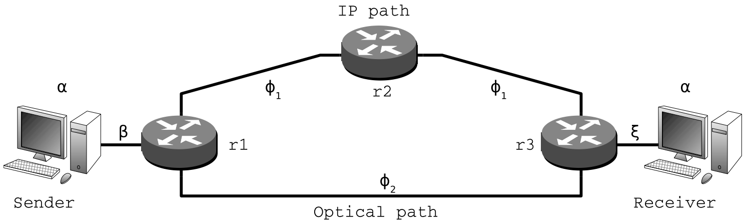

Fig. 4. Topology used in the simulations and limiting factors (Greek letters).

We observed different kinds of impact on the throughput of TCP flows. In all scenarios some throughput oscillation occurred during the transient phase, but TCP throughput recovered relatively fast after the transient phase was over. However, when the network link at the receiver side is the factor that limits the TCP throughput (thus the bandwidth of ξ is smaller than that of any other link), we found a huge impact on the performance of TCP. In this case, during the transient phase, router `r3` tries to send the last data received over the IP path, together with the first data received over the optical path, over the outgoing link ξ. The outgoing router's queue will be filled rapidly and packets must be dropped due to lack of queue space. It is interesting to note that the decrease in throughput was not caused by packet reordering, but by packet loss. This problem indicates that the transmission capacity of the link at the receiver side and the router's buffer size should be considered before moving flows on the fly.

7 Conclusions

Based on the research presented in this article, our main conclusion is that self-management is technically feasible to be deployed within hybrid optical and packet switching networks. From an implementation point of view, the decision process in a self-management approach can be built upon existing technologies, such as NetFlow/IPFIX to collect traffic information, and CLI, GMPLS, SNMP or NetConf to configure optical switches and routers.

Compared to traditional management approaches, self-management of hybrid networks provide several advantages. These advantages include better network performance, faster lightpath establishment and release, the ability to move large flows to the optical level, even in cases where such flows have not been made known to the human manager in advance, and finally the possibility to avoid congestion or underutilization by dynamically changing the decision thresholds.

Self-management of hybrid networks also has a number of disadvantages, however. An obvious disadvantage is that self-management increases the complexity of the management system, and therefore makes it harder to debug possible failures. Another problem is that large flows generated by a DoS attack can be moved to the optical level, and in this manner strengthen the attack. Also the move of an IP flow to the optical level may result in the temporary re-ordering and loss of packets, especially in cases where the bandwidth of the network links at the receiver side are the limiting factor that determine the TCP throughput.

Acknowledgments: This work has been supported by the EC IST-Emanics NoE (26854), as well as the EC IST-UniverSelf (FP7-257513) and the GigaPort Next Generation projects.

References

1. Asmare, E., Gopalan, A., Sloman, M., Dulay, N., Lupu, E.: A Mission Management Framework for Unmanned Autonomous Vehicles. In: Mobile Wireless Middleware, Operating Systems, and Applications, Second International Conference (Mobilware). ICST Lecture Notes, vol. 7, pp. 222–235 (April 2009)

2. Feldmann, A.: The Internet Architecture - Is a Redesign Needed?, pp. 147–164. Springer, Heidelberg (December 2009)
3. Fioreze, T.: Self-Management of Hybrid Optical and Packet Switching Networks. Ph.D. thesis, Universiteit Twente, Enschede (February 2010)
4. Fioreze, T., Granville, L., Pras, A., Sperotto, A., Sadre, R.: Self-management of hybrid networks: Can we trust netflow data? In: 11th IFIP/IEEE International Symposium on Integrated Network Management (IM 2009), Long Island, New York, USA. pp. 577–584 (June 2009)
5. Fioreze, T., Granville, L., Sadre, R., Pras, A.: A statistical analysis of network parameters for the self-management of lambda-connections. In: 3rd International Conference on Autonomous Infrastructure, Management and Security (AIMS 2009), Enschede, The Netherlands. vol. 5637, pp. 15–27 (July 2009)
6. Fioreze, T., van de Meent, R., Pras, A.: An Architecture for the Self-management of Lambda-Connections in Hybrid Networks. In: 13th EUNICE Open European Summer School, Enschede, The Netherlands. LNCS, vol. 4606, pp. 141–148. Springer Verlag, Germany (May 2007)
7. Fioreze, T., Pras, A.: Using Self-management for Establishing Light Paths in Optical Networks: an Overview. In: 12th EUNICE Open European Summer School, Stuttgart, Germany. pp. 17–20. Institut für Kommunikationsnetze und Rechnersysteme, Universität Stuttgart, Stuttgart, Germany (September 2006)
8. Fioreze, T., Pras, A.: Self-management of Lambda-connections in Optical Networks. In: Proceedings of the 1st International Conference on Autonomous Infrastructure, Management and Security (AIMS 2007) Student Workshop, Oslo, Norway. LNCS, vol. 4543, pp. 212–215. Springer Verlag, Berlin (June 2007)
9. Horn, P.: Autonomic computing: IBM's Perspective on the State of Information Technology (2001), `http://www.research.ibm.com/autonomic/manifesto/autonomic_computing.pdf`
10. Jennings, B., van der Meer, S., Balasubramaniam, S., Botvich, D., Foghlu, M., Donnelly, W., Strassner, J.: Towards autonomic management of communications networks. Communications Magazine, IEEE 45(10), 112 –121 (October 2007)
11. Leon-Garcia, A., Widjaja, I.: Communication Networks: Fundamental Concepts and Key Architectures. McGraw-Hill Companies, New York, 2nd edn. (2003)
12. Lupu, E., Dulay, N., Sloman, M., Sventek, J., Heeps, S., Strowes, S., Twidle, K., Keoh, S.L., Schaeffer-Filho, A.: AMUSE: autonomic management of ubiquitous e-Health systems. Concurrency and Computation: Practice and Experience 20(3), 277–295 (2008)
13. Miyazawa, M., Ogaki, K., Otani, T.: Multi-layer network management system with dynamic control of MPLS/GMPLS LSPs based on IP flows. In: The 11th IEEE/IFIP Network Operations and Management Symposium, 2008 (NOMS 2008), Salvador, Brazil. pp. 263–270 (April 2008)
14. Pras, A.: Network Management Architectures. Ph.D. thesis, Universiteit Twente, Enschede (February 1995)
15. Timmer, M., de Boer, P.T., Pras, A.: How to Identify the Speed Limiting Factor of a TCP Flow. In: 4th IEEE/IFIP Workshop on End-to-End Monitoring Techniques and Services (E2EMON2006), Vancouver, Canada. pp. 17–24 (April 2006)