Stability in dynamic networks

Stuart Clayman, Richard Clegg, Alex Galis and Antonio Manzalini[†]

Dept of Electronic Engineering, University College London, London, UK
Email: s.clayman@ucl.ac.uk, richard@richardclegg.org, a.galis@ucl.ac.uk

[†] Telecom Italia - Strategy / Future Centre, Turin Italy
Email: antonio.manzalini@telecomitalia.it

***Abstract*—This paper discusses the issues of stability within dynamic networks, and in particular we try to understand what stability is within this context. Traditionally networks have been viewed as being a relatively stable layer over which traffic is routed. The traffic flows and the routing updates have been seen as sources of instability. Recently, a large number of emerging and proposed networks have had the common characteristic of being dynamic in the sense that links and/or nodes can appear and disappear on a very short time scale (seconds or minutes). Such networks present a particular challenge in terms of the analysis of stability. In addition to the traditional methods by which instability can arise in a network, these networks have new modes by which instability can arise. This paper discusses these new forms of instability and potential issues thay may arise.**

I. Introduction

This paper discusses the subject of assessing stability in dynamic communications networks. The study of dynamic networks is a well-established field in mathematics which considers the property of time-varying graphs. In the case of this paper, the specific subject of study is computer networks which change their topology relatively rapidly in time. Traditionally networks have been viewed as being a relatively stable layer over which traffic is routed. Such networks add links and nodes relatively slowly as operators connect new machines to the network or add links between existing machines. However, an emerging class of networks has been identified which changes their behaviour rapidly in an automatic way. For example, in the case of a cloud network, new nodes are brought online when a task puts too much load on existing nodes. In the case of a mobile ad-hoc network, links between nodes are mediated by unreliable communication links which are a product of each node's position in space and these links change as the nodes move.

Within this paper stability is not to be understood in the formal technical sense embodied in mathematical formulations such as Lyapunov stability, Nyquist stability or Dijkstra's self-stabilisation concept for distributed algorithms [1]. Instead, stability and instability here simply refers to the notion of a system which exhibits "rapid" or "large" changes which may be detrimental to the system performance. The terms "rapid" and "large" are deliberately not defined here since these depend on context (a change in the overall topology of a network would be "rapid" if it took a few minutes but traffic on a link can change markedly in a sub-second timescale).

The paper presents a taxonomy of the areas of instabilities (in this looser sense) that will help in the categorization of instabilities in dynamic communications networks and how they can be assessed and addressed using various measures. In a traditional (non-dynamic) network, it is the traffic flows and the routing updates have been seen as sources of instability. Considerable progress has been made in understanding this area over the years, for an early review see [2]. For understanding stability in dynamic networks less progress has occurred, however [3] has addressed some issues and created some measures.

Recently, a large number of emerging and proposed networks have had the common characteristic of being dynamic in the sense that links and/or nodes can appear and disappear on a very short time scale (seconds or minutes). In addition to the traditional traffic flows and route changes by which instability can arise in a network, these networks have new modes by which instability can arise. Such networks present a particular challenge in terms of the analysis of stability. Initial attempts have been made in this area, including [4] and [5].

When stability is considered just over the domains of routing and the control plane, the problem has been researched for a number of years [6], [7]. Formal approaches to the problem include that of Kelly and Voice [8] who develop a decentralised algorithm for routing and rate control which is shown to be stable (in the sense that all points in the system converge to an optimal equibrium point) in a fluid flow approximation and of Hollot et al [9] who develop a dynamical systems approach to routing multipath TCP with guaranteed stability.

In this paper we discuss these new forms of instability and highlight the potential issues that may arise. Much of the work has been done within the context of the UniverSELF [10] project, which is investigating a unified management framework that targets the embedding of autonomic management in any type of network in a consistent manner. The paper has the following structure: Section II describes the different ways in which dynamic networks can be stable or non-stable. Section III describes some potential measures of instability. Section IV describes future research directions for this problem.

A. Types of dynamic communications network

The work presented here is within the field of highly dynamic networks: networks where nodes and links are regularly added or removed at short notice. This section provides a short summary of the dynamic networks context for the paper. Dynamic networks include virtual networks, logical networks, cloud computing networks, mobile ad-hoc networks, sensor networks and the Internet of Things.

In mobile ad-hoc networks (MANETS), links and nodes may appear and reappear spontaneously with no prior notice [11], [12]. Sensor networks are another environment where the dynamic network is extremely important [13], [14], [15]. In this context the limited power budget gives increased importance to reducing the overall network traffic. A common management approach for such networks is that data collection will occur at many nodes but data is sent to one of a set of chosen nodes (sometimes termed “cluster-heads") for aggregation. The problem is one of choosing cluster heads which minimise power drain but do not put much traffic on the network [16]. A recent, related idea is that of the Internet of Things [17], [18] where many millions of objects are used, for example via RFID tags, forming a highly dynamic network where nodes and links may appear and disappear quickly but the network is still existant.

A well-known dynamic network context is virtual networks [19], [20]. Virtual networks are a collection of virtual nodes connected together by a set of virtual links to form a virtual topology. In such networks, links and nodes may be reconfigured quickly and may be, for example, powered down to save energy or the node may be redeployed to a different logical area of the network. Both of these events are taken here to be equivalent to a node “death". On the other hand, virtual nodes may be brought online to deal with resources which are near their limits for bandwidth or CPU power. They are characterised in the literature either as a main means to test new Internet architectures or as a crucial component of future networks [21], [22].

Multiple logical networks can co-exist above the same physical substrate infrastructure. They can take the form of virtual private networks [23], programmable networks [24], overlay networks [19] or virtual networks [25]. The virtual nodes and links form a virtual topology over the underlying physical network.

An important area in dynamic networks are those aspects represented by the inter-connection and inter-operation of several heterogeneous dynamic networks sharing their resources, particularly in a virtualized manner. Resources such as processing, storage, and communication resources of multiple domains and networks can be made available for aggregation to support the provision of any service in a pervasive manner. Such networks are commonly called Software Driven / Software Defined Networks [26], [27], [28], [29].

Another area where highly dynamic networks are important is in the area of cloud computing. The EU project RESERVOIR [30], [31] studied federated cloud computing and the interactions between a distributed system of computing clouds. Each service in a federated cloud presents another type of dynamic network.

II. A TAXONOMY OF INSTABILITY

In this section a taxonomy is given for the areas of instability in dynamic communications networks. Throughout it should be remembered that this paper refers to stability in a less formal sense than concepts from dynamical systems.

In the case of traditional networks then instability can arise in various areas. The traffic itself can be unstable, perhaps because of large fluctuations in demand, perhaps because of protocol instabilities and so on. The management mechanisms for the traffic can produce fluctuating outputs, for example continual announcements and withdrawal of routes, which can be seen as an instability. The interaction between the varying traffic loads and the routing changes can lead to further instability. For example, the well-known phenomenon of route flap [32] occurs when a subset of prefixes has excessive numbers of announcements made about its reachability. This is, in itself, a management instability. However, if traffic is being routed to the prefixes in question then the management instability will lead to traffic instability as traffic to the affected prefixes rapidly switches its route through the network.

In addition to the management/control level and the traffic level in dynamic networks, further potential instability can be observed at the graph level. That is by the addition or removal of nodes and links which are part to the network’s structure. The very nature of a dynamic network means that such links and nodes will be added and removed on an ongoing basis as the network and the demands on it evolves. The question is whether such node or link changes will have an important effect on other network performance measures. For this reason the taxonomy will encapuslate the notion of *harmful* and *harmless* instabilty. Here instability is only considered harmful if it has a deleterious effect on the quality of service of an end user. For example, consider an ad-hoc network where the link to a particular node is extremely erratic and that link comes and goes repeatedly in a short timescale. Even if a network contained a large number of such nodes on its fringes, the overall impact on the network itself would be small. Conversely, a single node, well-placed in the heart of a dynamic network could create a large amount of disturbance if it semi-regularly disappeared and then reappeared long enough to insert itself back into well-used routes on the network. The notions of *harmful* and *harmless* effects are to some extent soft concepts and most forms of instabiltiy will form somewhere in between these extremes.

Figure 1 is a Venn diagram which shows the basis for our taxonomy of areas of network instability. Three areas where instability can occur are defined:

- The *graph* level refers to changes within the network when it is viewed as a graph in the mathematical sense.

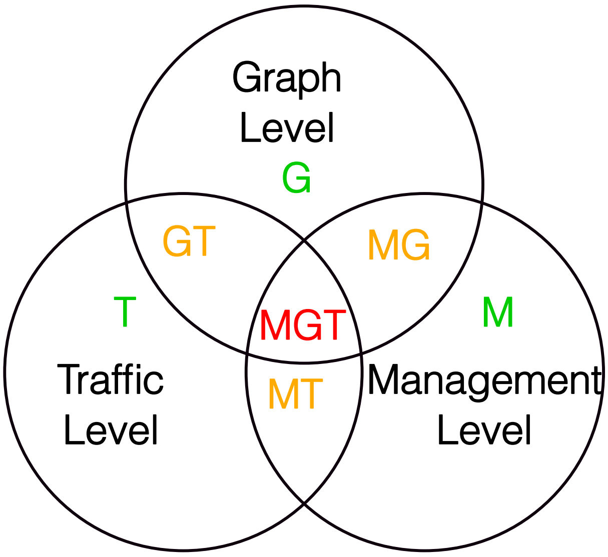

Fig. 1. Three area of network instability and their overlap

That is the addition or deletion of nodes and links.

- The *traffic* level refers to the changes of traffic level on one or more network links.
- The *management/control* level refers to changes in the management or control planes. For example, a change to routing or rules governing routing would be at this level.

The overlaps between circles indicate effects which arise jointly from the interaction at two or more of these levels so, for example, MG is an interaction between graph level and management/control level stability. It should be noted that the circles don't imply a direction of causality for the instability. So the MG (management/graph) part would include both examples where an unstable management system led to instability in the network graph but also examples where instability in the network graph led to an unstable management system. The colours in the diagram from green through amber to red indicate the growing complexity of the interaction.

G: Graph instability The graph level refers to the network when viewed as a graph, that is simply as nodes and links. Instability at this level would be the rapid addition or removal of nodes and links. Often such rapid additions or removals are a completely normal and non-harmful part of the operation of a dynamic communications network. For example, in an ad-hoc mobile network, it would be quite normal for a node with an unreliable connection to the rest of the network to add and remove itself from the network regularly as the link to its nearest node went up and down.

M: Management/control instability The management/control level refers to management plane and control plane instability within the network. The term control plane in networking originally applied to simply the part of the router architecture which was responsible for defining the onward direction of incoming packets. More recently the term control plane has been understood more broadly to include other functionality such as lightpath allocation, admission control and much else. For a definition of management plane see [33]. Here the management and control planes are considered together in their capacity as mechanisms which, ultimately, are responsible for layering a user demand for traffic onto the underlying network. A classic example of control plane instability would be the count-to-infinity problem in distance vector routing [34].

T: Traffic instability The traffic level refers simply to the measured traffic actually using a network link. At this level instability may represent varying user demand (a flashcrowd) or protocol interactions (such as TCP interactions). Instabilities at this level are well-studied as mentioned in the introduction.

MG: Management/graph instability This refers to interactions between the management and graph level which lead to instability. In this case the causality could go in either direction. For example, in a cloud network a poorly designed management plane could introduce instabilities at the graph level, perhaps by assigning nodes to work on an overloaded task and then removing them again as the task's loading dropped. Conversely, graph instability could cause more serious control plane instability. For example a poorly designed ad-hoc mobile network could introduce an unreliable node as a central part of its routing scheme. As this node entered and left the network there would be large scale routing instability as a result.

MT: Management/traffic instability As mentioned in the introduction, interactions between the control plane and network traffic are extremely well studied. The classic form of instability in this interaction is "route flap" where heavy traffic on a link causes a routing switch to a second link. This in turn causes heavy traffic on the second link and a routing switch back to the first. Less obvious forms of instability are possible. For example, a management system which had rules for "turn on traffic shaping if utilisation is above 80%" and "turn off traffic shaping if the usage is below 50KB/s" would be unstable for a link where a usage of 49KB/s represented utilisation of above 80%. While this is a very simple (and unlikely) example, much of the study of rule-based management systems is devoted to avoiding more indirect "loops" like this which could cause control instability.

GT: Graph/traffic instability Interactions between the graph level and traffic level could lead to a variety of stability problems. For example, a node which suddenly connects (or changes its connection point by swapping links) and places large demands for traffic on the network could overload individual links. Conversely, high loading on links could cause a node to become effectively cut off from the network.

MGT: Management/graph/traffic instability In this central section of the diagram is interactions between all levels. A large number of possibilities arise at this level depending on the exact nature of the network under study.

III. Measures of instability

In order to know whether there is instability in a network or if a particular suggested intervention on a network has improved stability we need to be able to measure stability itself. A large number of formal stability measures already exist, for example the Lyapunov exponent in chaotic systems. In the case of this paper, the measures are aimed at studying a more informal definition of stability. As before the measures can be aimed at the graph level, traffic level or management/control level. It is first necessary to look at how quantities change in time. All the measures chosen here will assume that the network is viewed at evenly spaced time intervals of length d – so the quantities measured are seen at time $0, d, 2d, \ldots$ and so on.

Given the observed time series, $f(t)$ measured at points $0, d, d2, \ldots$, then various measures of how much the time-series changes can be considered. The most obvious being variance, the mean absolute change between time steps and the range. All of these measures capture different aspects of stability. It is expected that the collected data from these time series can be visualized and analysed for any instability that arises.

In Figure 2 we illustrate how this could be presented. The figure presents and plots three time series (using artificial data to highlight our point) which have the same variance. *Time series 1 and 2* have the same range. *Time series 1* clearly has "rougher" behaviour and it might be considered more "unstable". From a network management perspective, each time series might be considered undesirable in different ways. *Time series 1* exhibits erratic change at each point in time (in fact it is simply generated from a normal distribution with unit variance). *Time series 2* (in fact generated simply by simply sorting time series 1) has the same mean and the same variance but very different behaviour. In some circumstances this time series might actually cause more problems for a network because it remains high for a long time. Finally *time series 3* has only two values, changes only once and has the same mean and variance as the other two. In some circumstances the single severe change might be considered extremely undesirable behaviour in a network but in most cases this would be easier behaviour for a network to handle. In addition to the variance, then, it will be useful to study the change between points, let this series be $f'(n) = f(n) - f(n-1)$.

A. Graph level stability measures

The upper layer at which stability could be assessed is to view the network as a graph. At this level the only changes which can be observed are those which affect the topology of the network, i.e., the addition or deletion of nodes and edges. In a traditional view of a wired network, the network graph is usually considered fixed, or so slowly changing as to affect little. However, in many more modern contexts (for example virtual networks, sensor networks, ad-hoc networks and mobile networks), nodes and edges can be short-lived. In this case we must consider graph stability.

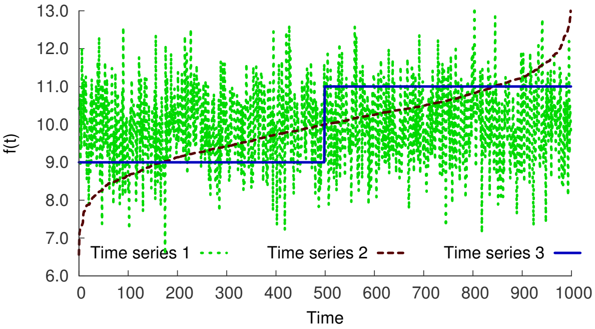

Fig. 2. Three time series illustrating variance and range

Let $G(t)$ be the graph at time t with $N(t)$ nodes and $E(t)$ edges at time t. In a real system, continuous observation is not possible so usually a time series with intervals is measured. From here and throughout the notation $G^{(n)}$ is used for the nth reading from $G(t)$ so $G^{(n)} = G(dn)$.

Some time series measures which might be observed on the graph $G(t)$ are:

- $N^{(n)}$ – Number of nodes at time nd.
- $E^{(n)}$ – Number of edges at time t.
- $d_{ij}^{(n)}$ – distance (in hops) from node i to j.
- $D^{(n)} \max_{i \neq j} d_{ij}(t)$ – graph diameter (distance between furthest separated points).
- $\overline{d^{(n)}} = \sum_{i \neq j} d_{ij}(t)/(N(dn)(N(dn)+1)$ – mean distance between nodes.

Many other measures could be considered, for example the betweenness-centrality for the different nodes on the network – this is a measure of how many shortest paths go through a node and large changes to this would indicate possible changes to a routing table. Yasinsac [3] has addressed some issues and has also created some measures.

B. Traffic level stability measures

The next level to consider are traffic related quantities within the graph. In this case observations may be made over a single or multiple links. Obvious quantities to observe for stability are:

- $B_i^{(n)}$ – bytes sent on link i in period $(d(n-1), dn]$.
- $P_i^{(n)}$ – packets sent on link i in period $(d(n-1), dn]$.
- $L_i^{(n)}$ – number of losses observed on link i in period $(d(n-1), dn]$.
- $R_{ij}^{(n)}$ – mean round-trip time from node i to node j in period $(d(n-1), dn]$.
- $\rho^{(n)}$ – utilization (the ratio of bytes sent in period $(d(n-1), dn]$ to the capacity of the link).

C. Management/Control plane level stability measures

Various indicators can be considered at the control plane level which could also be monitored for instability. In this case the quantities considered depend crucially on the management systems in place. Some examples might be: BGP updates, routing table updates and so on.

- $R^{(n)}$ – Number of routing updates at time nd.
- $B^{(n)}$ – Number of BGP updates at time nd.
- $C^{(n)}$ – Number of configuration commands sent to network devices at time nd.

Certainly more measures can be created for this area.

D. Viewing stability of nodes and links

Some of the quantities in the previous sections are properties of the network as a whole. For example, number of nodes and number of edges. Other quantities are properties of individual nodes and links. It would be, therefore, possible to take the measures associated with each node (or link) and create a vector representing all the quantities associate with that node (or link).

This allows introducing the concept of node, link or even sub-network state, which is a vector of quantities (data about relevant network parameters, e.g. QoS, etc) characterizing the state of node, link or even sub-network. Then a phase space (with dimensions of said vector) may represent the dynamic behavior in terms of states trajectories changing over time. The multi-variate time series resulting could then be assessed for stability in its various dimensions.

For practical reasons, the multi-dimensional phase space can be cut in a series of correlated 2D spaces and the evolution of a node could then be seen as a trajectory in 2D state space with desirable and undesirable regions. This phase space may have areas where an Operator wants the network state to be, and other areas where an Operator does not want the network to be. From this point of view, pursuing stability means avoiding abrupt phase changes in the phase space, specifically if they are moving network states into "not desired areas".

IV. Future directions

In this paper we have outlined the different kinds of dynamic networks that exist at present and highlighted the issue of dynamic network instability. It is clear that stability in dynamic networks will become an important topic, and we can expect that it will need to be addressed within network management systems in the future, and therefore some network management functions will need to be devised to measure and process instabilities in order to maintain networks and services.

We have presented a taxonomy of the areas of instabilities that will help in the categorization of instabilities in dynamic communications networks. Furthermore, we have shown some measures that can be assessed in order to determine if instabilities occur at the graph level, at the traffic level, and at the management/control plane level. This work is just a starting point for understanding this complex area.

As a consequence we will investigate the instability measures further and attempt to categorize further measures that can aid us. We intend to build a system that can monitor the highlighted dynamic networks and collect real data for the time series of the measures and determine what actually happens in such networks. From such data we will be able to make categorical statements about stability and instability, and allow network operators to control their networks in a more effective way.

Acknowledgment

This work is partially supported by the European Union UniverSELF project [10].

References

[1] E. W. Dijkstra, "Self-stabilizing systems in spite of distributed control," *Communications of the ACM*, vol. 17, no. 11, pp. 643–âĂŞ644, November 1974.

[2] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhauser, 2003.

[3] A. Yasinsac, "Rates of change in ad hoc networks," pp. 166–184, 2005.

[4] F. Kuipers, H. Wang, and P. Van Mieghem, "The stability of paths in a dynamic network," in *Proceedings of the 2005 ACM conference on Emerging network experiment and technology*, ser. CoNEXT '05. New York, NY, USA: ACM, 2005, pp. 105–114. [Online]. Available: http://doi.acm.org/10.1145/1095921.1095936

[5] F. Theoleyre and F. Valois, "About the self-stabilization of a virtual topology for self-organization in ad hoc networks," in *Proceedings of the 7th international conference on Self-Stabilizing Systems*, ser. SSS05. Berlin, Heidelberg: Springer-Verlag, 2005, pp. 214–228. [Online]. Available: http://dx.doi.org/10.1007/11577327_15

[6] R. Goodman and B. Ambrose, "Stability of traffic patterns in broadband," *Journal of Network and Systems Management, Special Issue on Routing in Broadband Networks*, vol. 3, no. 4, pp. 371–380, December 1995.

[7] M. A. Marsan, M. Franceschinis, P. Giaccone, E. Leonardi, F. Neri, and A. Tarello, "Instability phenomena in underloaded packet networks with elastic traffic," 2003.

[8] F. Kelly and T. Voice, "Stability of end-to-end algorithms for joint routing and rate control," *Computer Communication Review*, vol. 35, no. 2, pp. 5–12, 2005.

[9] H. S. Hollot, H. H. S. Shakkottai, C. V. Hollot, R. Srikant, and D. Towsley, "Overlay tcp for multi-path routing and congestion control," in *In ENS-INRIA ARC-TCP Workshop*, 2004.

[10] UniverSELF consortium, "UniverSELF project," http://www.univerself-project.eu/.

[11] A. Malatras, G. Pavlou, and S. Sivavakeesar, "A programmable framework for the deployment of services and protocols in mobile ad hoc networks," *IEEE Transactions on Network and Service Management*, vol. 4, no. 3, pp. 12–24, Dec. 2007.

[12] S. S. Yau, F. Karim, Y. Wang, B. Wang, and S. K. S. Gupta, "Reconfigurable context-sensitive middleware for pervasive computing," *IEEE Pervasive Computing*, vol. 1, no. 3, pp. 33–40, 2002.

[13] C. Olston, B. T. Loo, and J. Widom, "Adaptive precision setting for cached approximate values," in *ACM SIGMOD*. New York, NY, USA: ACM, 2001, pp. 355–366.

[14] G. Cormode, M. Garofalakis, S. Muthukrishnan, and R. Rastogi, "Holistic aggregates in a networked world: distributed tracking of approximate quantiles," in *ACM SIGMOD*, 2005, pp. 25–36.

[15] A. Sharaf, J. Beaver, A. Labrinidis, and K. Chrysanthis, "Balancing energy efficiency and quality of aggregate data in sensor networks," *The VLDB Journal*, vol. 13, no. 4, pp. 384–403, 2004.

[16] K. Akkaya, F. Senel, and B. McLaughlan, "Clustering of wireless sensor and actor networks based on sensor distribution and connectivity," *J. Parallel Distrib. Comput.*, vol. 69, pp. 573–587, June 2009.

[17] E. Welbourne, L. Battle, G. Cole, K. Gould, K. Rector, S. Raymer, M. Balazinska, and G. Borriello, "Building the internet of things using RFID: The RFID ecosystem experience," *Internet Computing, IEEE*, vol. 13, no. 3, pp. 48–55, 2009.

[18] L. Atzori, A. Iera, and G. Morabito, "The internet of things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[19] N. M. K. Chowdhury and R. Boutaba, "Network virtualization: State of the art and research challenges," *IEEE Communications Magazine*, vol. 47, no. 7, 2009.

[20] M. Casado, T. Koponen, R. Ramanathan, and S. Shenker, "Virtualizing the network forwarding plane," in *PRESTO, CONEXT Workshop*, 2010.

[21] T. Anderson, L. Peterson, S. Shenker, and J. Turner, "Overcoming the internet impasse through virtualization," *Computer*, vol. 38, pp. 34–41, April 2005.

[22] N. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Computer Networks*, vol. 54, no. 5, pp. 862 – 876, 2010.

[23] L. Andersson and T. Madsen, "Provider provisioned virtual private network (VPN) terminology," *Internet Engineering Task Force, RFC 4026*, March 2005.

[24] A. Galis, S. Denazis, C. Brou, and C. Klein, *Programmable Networks for IP Service Deployment*. Artech House Books, 2004.

[25] T. Anderson, L. Peterson, S. Shenker, and J. Turner, "Overcoming the internet impasse through virtualization," *Computer*, vol. 38, pp. 34–41, April 2005.

[26] "Open Networking Foundation," http://www.opennetworking.org/, 2012.

[27] "Open Flow," http://www.openflow.org/, 2012.

[28] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008. [Online]. Available: http://doi.acm.org/10.1145/1355734.1355746

[29] J. Rubio-Loyola, A. Galis, A. Astorga, J. Serrat, L. Lefevre, A. Fischer, A. Paler, and H. Meer, "Scalable service deployment on software-defined networks," *Communications Magazine, IEEE*, vol. 49, no. 12, pp. 84 – 93, december 2011.

[30] B. Rochwerger, D. Breitgand, D. Hadas, I. Llorente, R. Montero, P. Massonet, E. Levy, A. Galis, M. Villari, Y. Wolfsthal, E. Elmroth, J. Caceres, C. Vazquez, and J. Tordsson, "An architecture for federated cloud computing," *Cloud Computing*, 2010.

[31] B. Rochwerger, D. Breitgand, E. Levy, A. Galis *et al.*, "The RESERVOIR model and architecture for open federated cloud computing," *IBM Journal of Research and Development*, vol. 53, no. 4, 2009.

[32] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," RFC 2439 (Proposed Standard), Internet Engineering Task Force, Nov. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2439.txt

[33] D. Macedo, Z. Movahedi, J. Rubio-Loyola, A. Astorga, G. Koumoutsos, and G. Pujolle, "The AutoI approach for the orchestration of autonomic networks," *Annals of Telecommunications*, vol. 66, pp. 243–255, 2011.

[34] G. Malkin, "RIP Version 2," RFC 2453 (Standard), Internet Engineering Task Force, Nov. 1998, updated by RFC 4822. [Online]. Available: http://www.ietf.org/rfc/rfc2453.txt