ABSTRACT

In recent years, we witnessed the evolution of the Internet from a simple means of exchanging information in a timely and cost effective manner; to an instrumental tool in service delivery. To this date, the Internet is perceived as possessing tremendous potential in socio-economic development due to the wide range of applications and services it has to offer. These applications are increasingly becoming both bandwidth and delay intensive as in the case of Telemedicine, Video-Conferencing and e-learning. Moreover, business reliance on the Internet is posing more stringent requirements on availability, reliability and Quality of Service (QoS), especially in e-commerce applications.

Many techniques aim to enhance the efficiency and robustness of the Internet by granting end hosts access to multiple paths through the Internet, and enabling them to control which traffic traverses each path. This results in improved security, resilience, and load balancing. One such technique is multihoming, where an enterprise has multiple links to the Internet from different Internet Service Providers (ISPs) used for redundancy and traffic distribution. Multihoming is typically implemented using Border Gateway Protocol (BGP) routing protocol [RFC1771] which maintains a table of IP networks or 'prefixes' designating network reachability among Autonomous Systems (AS). BGP is implemented on the edge routers of the ASs and chooses the shortest route based only on the AS hop counts as well as the policies set by each AS. Currently, multihoming and BGP routing are used by ISPs and large enterprises not only to increased reliability and redundancy, but also to meet their SLAs by reducing delay and enforcing a fault management system.

Very few authors focused on analyzing network architectures in terms of path failure probabilities, which reduces the number of available paths, increases load on remaining paths

and affects performance. In reference [40] a model for multihomed, BGP routing networks was analyzed with a focus on reliability; defined as the probability of continuing transmissions on link failures. This model, though apt for the stated purpose, fails to capture the side effects of partial network failures on meeting a specific SLA objective.

The focus of this thesis is on finding the probability of all possible network states, which is defined as the state of all links in the network. Subsequently, the use of these probabilities to quantify the effects of failures on network performance is demonstrated. A generic and flexible network structure is used, based on that used in [40]. However, some modifications to the network parameters are made to enable description of more practical structures. In particular, the structure is divided into enterprise, ISP and SOHO sections, but unlike [40] the link failure rates as well as the number of multipath within the ISP sections can be varied independently.