Research Statement for Theophilus Benson

The average enterprise network operator is responsible for managing a network that secures $5 billion worth of intellectual property (McAfee 2010) and whose poor performance can result in the loss of $100 million in operating revenue (InformationWeek 2007). My current and future research are motivated by the importance of such networks and focus on developing fundamental models that make them easier to understand, and on designing frameworks that make them more secure, more reliable and more efficient. My research has had a broad impact on the theory and practice of networking; my work on modeling the complexity of network configurations has helped to discover latent security holes and showed how to systematically mitigate future misconfiguration; my work on data centers has resulted in crucial insights into existing data center traffic behavior and performance issues; and my work on cloud computing has been adopted by one of IBM's research groups.

I have accomplished this by uniquely viewing the network not as an isolated system of devices and links to be studied, designed, optimized and tweaked, but as a sub-component whose performance, reliability, and security must be viewed in the larger context of the systems interacting with it. This view naturally complicates the problem but also opens up the solution space and creates opportunities for the integration of techniques and algorithms from other areas of computer science, such as, programming languages and software engineering.

CURRENT RESEARCH: MANAGING ENTERPRISE NETWORKS

My current work focuses on *improving the security, reliability and efficiency of enterprise networks* by systematically analyzing the fundamental choices available in designing and configuring them. The challenge lies in dealing with device heterogeneity, with a plethora of low-level options, and with a multitude of complementing and interacting protocols. My work views an enterprise network as a holistic system comprising of the campus network interconnecting employees, the ISP services interconnecting different enterprise campuses, the data centers and clouds housing most of the processing and storage units, and it even includes the operators configuring these networks. My work has resulted in the design of complexity metrics, a novel framework to reason about network design; the development of policy units, an application that informs the design of enterprise networks; and the first-ever analysis of the design of the ISP services used by enterprise networks.

Network Configuration Management: I argue that the most direct route to mitigating misconfiguration and improving security and reliability is to first understand the factors that introduce network complexity and then use these factors as guiding principles to control, and even eliminate, complexity. Although anecdotal evidence shows that design complexity is the cause of network misconfiguration and ultimately results in security and reliability issues in enterprise networks, the community has focused on fixing the symptoms (e.g., developing rules that automatically correct known misconfigurations) rather than addressing the root-cause of misconfiguration.

To better understand the factors that impact misconfiguration, I developed a family of *complexity metrics* (NSDI 2009) that abstracts away configuration details and succinctly describe the underlying complexity of a network's design. Inspired by programming languages and software engineering, I have designed my metrics to view the network as a distributed program, extracting dependencies between different components, capturing uniformity, and modeling routing protocols to determine logical relationships. An interesting

observation from my work is that much of the complexity is added by non-technical factors, such as, financial constraints and locally-established common practices. This deeper understanding of network complexity and misconfiguration has provided me with the insights necessary to develop new applications that simplify current networks. For example, policy units (IMC 2009) automatically captures and presents the global reachability policies which are usually undocumented and buried within network configurations. Thus, automatically deriving and succinctly representing these policies streamlines the task of making changes and debugging the network. I have shown the effectiveness of my applications by applying them to 10 enterprise networks and unearthing configuration bugs that would have resulted in security holes or black holes. For example, in one network we discovered a misconfiguration of the routing protocol that would have disconnected an entire department. Another interesting consequence of my models and metrics is that they provide a set of abstractions from which we can design configuration interfaces for future platforms such as Software Defined Networks.

Although my metrics and applications apply to an enterprise's campus network, end-to-end performance and reliability are also affected by the ISP services providing connectivity to different enterprise campuses. To this end, I have conducted the first-ever methodical study of the complexity underlying the configuration of these enterprise centric ISP services. Using data from the AT&T service backbone, I found that complexity, and ultimately an enterprise's performance, can be critically impacted by seemingly mundane trade-offs, such as, the choice of vendor and the choice of service-specific routing design. My study (SIGCOMM 2011) provided a system that is the first step towards a new solution space: systematically navigating the spectrum of trade-offs to simplify network service design, thereby supporting improved enterprise performance.

Data Center Networks: Enterprises heavily employ data centers and clouds to host a variety of business-critical processing and storage units. As such, the productivity of an enterprise is heavily dependent on the performance of its data centers. My work on improving data center performance, and ultimately enterprise productivity, has lead to key insights on the performance of data center data planes and on technologies to enhance them. I have conducted a study on the data plane of data centers, developed a traffic engineering technique to improve its performance, and designed a more expressive virtualization layer for the data plane that encompasses both the physical network and servers.

My study (IMC 2010) on traffic characteristics of data centers is the largest-ever and most influential study of its kind: I have made wide reaching discoveries ranging from data center performance attributes, such as the fact that most losses occur at the edge of the network, to flow level properties, such as the fact that traffic is bursty within data centers. These findings are currently being used by researchers as guiding principles in designing their own data center frameworks and this work (IMC2010, WREN2009) has received a total of 74 citations in a little over a year. An interesting observation from this study is that the current data plane can effectively support existing traffic demands given appropriate network mechanisms and that the data plane can be controlled in a centralized fashion. Motivated by these findings, I have developed a novel traffic engineering technique (CoNEXT 2011) to improve the performance of existing data centers without forcing enterprises to undergo expensive upgrades to their data center networks. This is a crucial advancement as it helps postpone the adoption of new architectures, thus giving researchers time to develop techniques that are better suited for the traffic characteristics of data centers.

To further improve productivity and utilization, some enterprises have started to virtualize the data plane of their data centers and transform them into clouds; however, many are prevented by the cloud's lack of support for security, performance and reliability features. To address this shortcoming, I have developed CloudNaaS (SOCC 2011) a system that leverages indirection and the fine-grained control provided by programmable

devices to provide a more expressive virtualized data plane. This allows enterprises to easily replicate their current arbitrarily complex data plane configuration in the cloud. I believe that one metric for success as a researcher is that your results be not only applicable to real world systems but applied to them. My prototype of CloudNaaS exemplifies such success; it has been adopted by IBM research as the basis for ongoing cloud networking projects and, more recently, aspects of it have been licensed by a large cloud provider.

FUTURE RESEARCH AGENDA

In the future, I intend to continue applying my interdisciplinary and holistic approach towards research to develop more broadly applicable and unifying frameworks and models for network design and management.

Towards a Grand Theory of Network Design: My current complexity metrics focus on the syntax used to express network policies. However, an important observation of my work is that much of the syntactic complexity in modern systems is added by non-technical factors, such as, financial constraints and locally-established common practices. Two interesting research directions arise from this. The first involves developing a set of models to capture semantic complexity or the additional complexity introduced by the choices operators make when selecting the protocols and primitives to use to implement an enterprise's policies. For example, to restrict access to confidential HR resources, one operator may use simple data plane primitives while another may decide to implement the exact same policy by modifying control plane protocols. While both approaches achieve the same objective, the higher complexity of implementing the latter increases the likelihood of misconfiguration. Second, given that the non-technical factors discussed earlier are likely to play a key role in the future, even as new languages and abstractions are introduced, e.g. in the context of Software Defined Networking, I believe that it is crucial for network researchers and network operators to understand the complexity of representing network policies and other non-technical factors in future configuration languages. To this end, I intend to generalize my current complexity metrics to quantify both the syntactic and semantic complexities in future and higher layer configuration languages. This requires new models for capturing dependencies, uniformity, and logical relationships between components and ways to compare two different practices in an apples-to-apples fashion.

Application-Aware Data Centers: My current work on data centers has been largely agnostic of the applications running within the data center. However, I believe that much greater performance can be extracted from the data plane when both the applications and the network cooperate with each other. For example, hints from an application about its intended communication pattern can be used to smoothen out the burstiness of traffic and ultimately reduce congestion and packet losses. Considering applications in the design of data center networks and protocols introduces new research questions. For example, how do you build a specialized data center with custom servers and network switches that are optimized for a given class of applications? How do you design protocols that enable better statistical multiplexing by utilizing intimate knowledge about the different applications running within the data centers? How do you facilitate sharing of hints for optimizing network performance, such as application communication patterns and behavior, between a cloud tenant and cloud provider while ensuring that both privacy and security are preserved? Although initial steps are being taken, by entities such as YouTube and Bing to develop application-specific data centers, little work has been done to develop a systematic and generalizable method for creating such data centers for more complex and intricate applications.