Chapter 10
A Top–Down Framework for Modeling Routing Design Complexity

Xin Sun
Florida International University, USA

ABSTRACT

Networks with higher degrees of complexity typically require more effort to manage and are more prone to configuration errors. Unfortunately, complexity remains one of the least understood aspects of networking. This chapter takes a first step toward bridging this gap, by presenting a top-down framework for modeling complexity in routing design, a critical and highly complex network design task. First, a set of abstractions is introduced for precisely defining objectives of routing design, and for reasoning about how a combination of routing design primitives will meet the objectives. Next, models are presented for quantitatively measuring the complexity of a routing design by modeling individual design primitives and leveraging configuration-derived complexity metrics. This modeling approach helps understand how individual design choices may impact the resulting complexity, and enables comparison of alternative routing designs and "what-if" analysis of the potential impact of a design change on complexity.

INTRODUCTION

Ad-hoc design decisions made during organic growth (e.g., company mergers, expansion to new markets) and incremental evolution (e.g., user mobility, department reorganizations) have caused many enterprise networks to become unnecessarily complex and increasingly difficult and costly to operate. Operator interviews and anecdotal evidence suggest that networks with higher degrees of complexity in their design and

implementation generally require more manual intervention to manage, are more difficult to reason about, predict and troubleshoot, and are more prone to configuration errors. For many complex enterprise networks, the amount of management effort required has become the dominant cost of operation (Kerravala, 2004). Despite this investment, configuration errors account for 50%-80% of network outages (Juniper, 2008), 80% of Air Force network vulnerabilities (CSIS, 2008), and 65% of all successful cyber-attacks (Pescatore,

2003). Further, the significant amount of time, effort and risk involved in upgrading complex networks has forced many operators to be extremely reluctant to adopting new and innovative architectures (e.g., software defined networking, cloud computing), and thus poses a fundamental barrier to innovation in the long term.

Unlike network performance, resiliency and security which have all been extensively studied and for which models, algorithms and tools have been successfully developed, complexity remains the least understood part of computer networks. While part of the complexity in network design is inherent, given the wide range of operational objectives that these networks must support, to include security (e.g., implementing a subnet level reachability matrix), resiliency (e.g., tolerating up to two component failures), safety (e.g., free of forwarding loops), performance, and manageability. There is also evidence, however, to suggest that some of the network design complexity may have resulted from a semantic gap between the high-level design objectives and the diverse set of routing protocols and low-level router primitives for the operators to choose from. Often, for the same target network multiple designs exist to meet the same operational objectives, and some designs are significantly easier to implement and manage than others. For example in some cases, route redistribution may be a simpler alternative to BGP (Border Gateway Protocol) for connecting multiple routing domains. However, the reasoning about complexity remains qualitative and objective, and essentially a black-art. Today we do not have a scientific model to characterize the degree of complexity in a given network design, or to compare two design proposals in terms of their complexity. Lacking an analytical model to guide the operators, the current routing design process is mostly ad hoc, prone to creating designs more complex than necessary. This lack of research on network complexity has been more recently recognized by both academia and industry.

As a first and important step towards bridging this gap, this chapter seeks to quantitatively model the complexity associated with a network design, with a view to developing alternate designs that are less complex but meet the same set of operational objectives. In doing so, it focuses on routing design as a concrete application, given its prevalence and high complexity. More specifically, the chapter presents a top-down approach to characterizing the complexity of enterprise routing design given only key high-level design parameters, and in the absence of actual configuration files. The model takes as input abstractions of high-level design objectives such as network topology, reachability matrix (which pairs of subnets can communicate), and design parameters such as the routing instances, and choice of connection primitive (e.g., static routes, redistribution, etc.).

The overall modeling approach is to (i) formally abstract the operational objectives related to the routing design which can help reason about whether and how a combination of design primitives will meet the objectives; and (ii) decompose routing design into its constituent primitives, and quantify the configuration complexity of individual design primitives using the existing bottom-up complexity metrics.

A top-down approach such as this has several advantages. By working with design primitives directly (independent of router configuration files), the model is useful not only for analyzing an existing network, but also for "what if" scenarios analysis capable of optimizing the design of a new network and similarly, a network migration, or evaluating the potential impact of a change to network design. Further, the presented models help provide a conceptual framework to understand the underlying factors that contribute to configuration complexity. For example, reachability restrictions between subnet pairs may require route filters or static routes, which in turn manifest as dependencies in network configuration files.

Related Content

A Culture of Survivors: SlutWalk, Third Culture, and New Media Communication
Jennifer L. Seifert (2015). *International Journal of Interactive Communication Systems and Technologies (pp. 1-13).*
www.irma-international.org/article/a-culture-of-survivors/136848

Interactive TV Together: An Open Service Infrastructure for Enhancing Interactive TV Experiences
Cristian Hesselman, Joost Broekens, Mark Gülbahar, Florian Winkler, Daniel Görgenand Ferry de Jong (2009). *Social Interactive Television: Immersive Shared Experiences and Perspectives (pp. 30-49).*
www.irma-international.org/chapter/interactive-together-open-service-infrastructure/29198

Digital Activism and Digital Revolution in Objective Journalism
Ugochukwu Matthew Okwudiliand Jazuli S. Kazaure (2020). *International Journal of Interactive Communication Systems and Technologies (pp. 39-56).*
www.irma-international.org/article/digital-activism-and-digital-revolution-in-objective-journalism/275195

A Forecast of the Adoption of Wearable Technology
Tom Page (2018). *Wearable Technologies: Concepts, Methodologies, Tools, and Applications (pp. 1370-1387).*
www.irma-international.org/chapter/a-forecast-of-the-adoption-of-wearable-technology/202018

Enhancing Self-Reflection With Wearable Sensors Workshop: A Commentary on the ACM MobileHCI 2014 Workshop
Genovefa Kefalidou, Vicky Shipp, James Pinchinand Alan Dix (2018). *Wearable Technologies: Concepts, Methodologies, Tools, and Applications (pp. 628-636).*
www.irma-international.org/chapter/enhancing-self-reflection-with-wearable-sensors-workshop/201978