

An Integrated Systematic Approach to Designing Enterprise Access Control

Xin Sun and Geoffrey G. Xie

Abstract—Today, the network design process remains *ad hoc* and largely complexity agnostic, often resulting in suboptimal networks characterized by excessive amounts of dependence and commands in device configurations. The unnecessary high configuration complexity can lead to a huge increase in both the amount of manual intervention required for managing the network and the likelihood of configuration errors, and thus must be avoided. In this paper, we present an integrated top-down design approach and show how it can minimize the unnecessary configuration complexity in realizing reachability-based access control, a key network design objective that involves designing three distinct network elements: virtual local-area network (VLAN), IP address, and packet filter. Capitalizing on newly developed abstractions, our approach integrates the design of these three elements into a unified framework by systematically modeling how the design of one element may impact the complexity of other elements. Our approach goes substantially beyond the current divide-and-conquer approach that designs each element in complete isolation, and enables minimizing the combined complexity of all elements. Specifically, two new optimization problems are formulated, and novel algorithms and heuristics are developed to solve the formulated problems. Evaluation on a large campus network shows that our approach can effectively reduce the packet filter complexity and VLAN trunking complexity by more than 85% and 70%, respectively, when compared with the *ad hoc* approach currently used by the operators.

Index Terms—Network management, network complexity, top-down network design, access control, VLAN, IP addressing, packet filters.

I. INTRODUCTION

RECENT research [7], [23], [33] and vendor documents [1], [30] reveal that multiple, distinct routing designs are possible to meet the same set of enterprise network operational requirements (e.g., security policy represented by a reachability matrix [38]). Moreover, the configuration complexity of these designs can vary greatly. In other words, some designs may incur much higher configuration complexity than others while accomplishing the same objectives. The *unnecessarily high* configuration complexity is highly undesirable as it can lead to a huge increase in both the amount of

manual intervention required for managing the network and the likelihood of configuration errors. For example, a research report [21] discloses that 80% of enterprise IT budget is devoted to maintaining the status quo. Despite this investment, configuration errors account for 50-80% of network outages [19], [21] and enable 65% of all successful cyber-attacks [31]. There is a general perception that complexity is the primary cause of high human costs, and interviews and anecdotal evidence suggest that an operator's ability to run a network decreases as the network becomes more complex [7].

Thus, an important open research question arises: *Is it possible to systematically identify, among all designs that can meet given operational requirements, the one(s) with the minimum amount of configuration complexity?*

The current state of network design practice by operators is mostly *ad hoc* and, in particular, does not rigorously formulate the goal of minimizing network complexity. As a result, a large number of existing production networks may not be optimal in terms of configuration complexity [23], [35], likely causing a huge increase in operational costs. Having recognized the importance of the problem and associated challenges, researchers have recently begun to investigate this problem in the specific context of enterprise network design [34], [35]. These approaches focus primarily on meeting the specific objective of *user reachability control* (essentially implementing a subnet-level reachability matrix). They enable an operator to formulate an individual design task, such as grouping hosts of his/her network into different VLANs, into a model of optimizing a desired performance metric subject to a set of correctness and feasibility constraints.

While these recent advances in systematic network design create a major opportunity to address the complexity problem, the current approaches suffer from a critical limitation: they employ an oversimplified “divide-and-conquer” (i.e., stage-by-stage) strategy that models individual design steps in complete isolation even when the steps together implement a common goal. For example, totally independent formulations and optimality criteria are used for VLAN design and packet filter design [35] even though the two design steps share a common objective of user reachability control. While these formulations can potentially minimize the complexity of configuration at each design stage in isolation, the overall complexity may still be unnecessarily high. This is because the design choices made at an early stage (e.g., VLAN design or IP address allocation) can significantly affect the available design space of a later stage (e.g., packet filters), potentially resulting in a substantial amount of unnecessary complexity.

Manuscript received February 24, 2015; revised September 13, 2015; accepted January 25, 2016; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Sen. This work was supported by the National Science Foundation under Grant CRII CNS-1459761.

X. Sun is with the School of Computing and Information Sciences, Florida International University, Miami, FL 33199 USA (e-mail: xinsun@cs.fiu.edu).

G. G. Xie is with the Department of Computer Science, Naval Postgraduate School, Monterey, CA 93943 USA (e-mail: xie@nps.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2016.2535468

In this paper, we investigate a novel integrated top-down methodology that jointly designs multiple networking elements involved in achieving a common objective. As a first step in this important direction, we focus on designing new (i.e., “green-field”) networks.¹ The key components of our approach include: (i) for a given design objective, identifying all the networking elements that may be involved in its implementation, and their interactions (i.e., how the design of one element could affect the design of others); (ii) characterizing the source of complexity in each element, leveraging recently-developed complexity metrics; (iii) formulating the design problem as one of minimizing the total complexity of all involved elements, subject to correctness and feasibility constraints; and (iv) developing specific algorithms and heuristics to solve the formulated problems. As such, this new approach goes substantially beyond the state-of-the-art “divide-and-conquer” approaches. It requires not only entirely new formulations and algorithms, but also fundamentally new abstractions and models in order to integrate the design of multiple network elements into a unified framework.

Our integrated design methodology is general and can be applied to a variety of network design objectives and scenarios. In order to demonstrate its feasibility and power at sufficient depths, in this paper we focus on one concrete application: reachability-based access control. We choose access control because security is of vital importance to virtually every enterprise network, and also because its design involves multiple networking elements and as such it is highly challenging and will benefit greatly from the new approach.

Similarly, while our approach is agnostic to the type of network complexity metric used,² the focus of this paper is on minimizing the *configuration complexity*, specifically the amount of *command dependencies* [7] in the router configurations of the resulting network. According to recent studies [7], [33], [34], these dependencies are directly linked to the operational cost as they require substantial manual effort to configure correctly in the initial implementation and manage in subsequent evolutions, and if not maintained properly, can lead to serious issues such as application performance degradation and security breaches.

We evaluate the benefits of the new approach in the context of the heuristics we have developed for solving the formulated design problem of user reachability control. The evaluation is conducted on a large university campus network with several thousand user hosts. The results show that our approach can effectively reduce the number of packet filter rules and the number of VLAN trunk ports by more than 85% and 70%, respectively, when compared to the ad-hoc approach currently used by the operators.

The rest of the paper is organized as follows. In Section II, we briefly survey the state of the art. In Section III, we first substantiate the need for the integrated design approach using a detailed example of reachability control design.

¹Section VII provides a brief discussion of possible extension of this work to existing “brown-field” networks.

²Section VII provides a brief discussion of other potential complexity metrics.

We then demonstrate how the integrated approach can be applied to reachability-based access control, and formally introduce an integrated design framework for this problem. In Sections IV and V, we present new formulations and novel heuristics to accomplish the two design problems identified by the framework: joint design of VLAN and packet filters, and joint design of IP allocation and packet filters. Section VI presents a thorough evaluation of our heuristics in two campus network settings. Possible extensions and open issues are discussed in Section VII. Finally, we conclude the paper and briefly outline our plan for future work in Section VIII.

II. STATE OF ART OF NETWORK DESIGN

In this section, we overview the current state of the art of enterprise network design, specifically focusing on the more recent developments in top-down design techniques. Our aim is to not only discuss related work, but also provide a historical perspective of the proposed integrated design approach before we present detailed examples to substantiate how the new approach may reduce complexity in the next section.

A. Operational Practice and Tools

The operational community has a rich history of crafting the art of network design and reconfiguration. Nonetheless, the state of the practice by operators is still defined predominantly by ad-hoc, manual decision making. Notable efforts to simplify network design involve template-based approaches that codify and promote best practices [1]–[4] and abstract languages to specify configurations in a vendor-neutral fashion [12]. There are also tools such as PRESTO [13] to convert a network design into device-vendor-specific configuration commands. These approaches merely model the low-level mechanisms and their configuration. They do not model network-wide operator intent such as reachability and manageability. A logic-based approach to configuration generation based on model-finding is presented in [28]. The focus is on the generation of configuration parameters conforming to correctness rules distilled from best practices, and the system does not take complexity into consideration. Many works have approached the problem of minimizing the number of rules in a *single* packet filter (e.g., [26]). In contrast, we focus on minimizing the total number of filter rules required for a given network to meet all its access control requirements.

Finally, various design guidelines including those for a top-down network design approach [30] can be found in the literature. These guidelines provide practical insights into the trade-offs of different design choices regarding topology, hardware and protocols. However, considerable manual effort is required to determine how to apply these guidelines to the design of a network of medium to large size.

B. Systematic Multi-Stage Design

Systematic network design, characterized by the use of a formal model to generate configuration that is *provably* correct and additionally *optimizes* certain performance metrics, has emerged as a potential solution to the challenges facing the operational community. Early efforts on this front focus

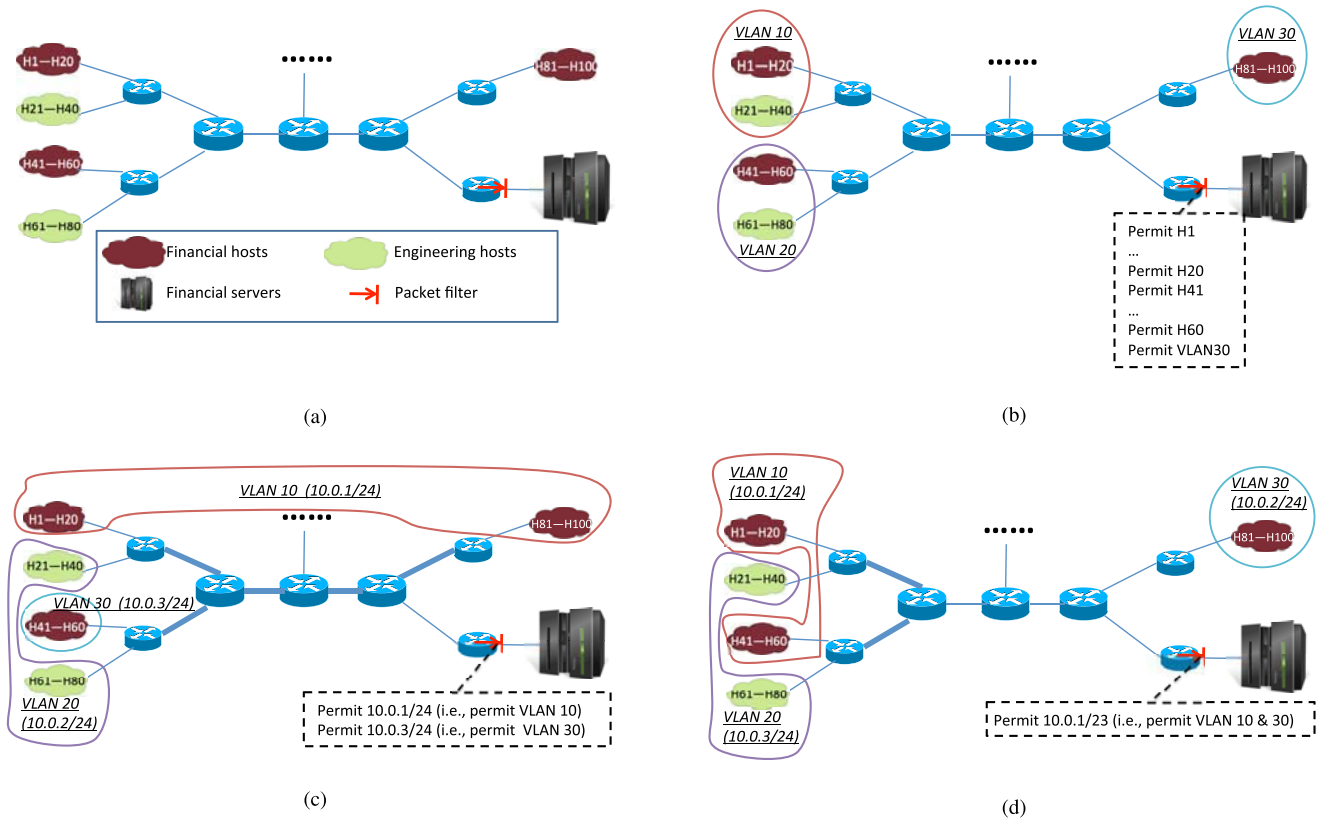


Fig. 1. Multiple designs with different complexity characteristics exist for one network. (a) The example network to be designed. (b) Design #1: Purely location-based VLAN grouping leads to an explosion in the number of filter rules. (c) Design #2: This VLAN grouping leads to fewer filter rules but an excessive number of trunk ports. Also, the IP allocation scheme does not allow aggregation of filter rules. (d) Design #3: A better VLAN grouping scheme that reduces the sum of filter rules and trunk ports. Also a better IP allocation scheme that facilitates aggregation of filter rules.

on tasks encountered in carrier networks, such as configuring BGP policies [9], [14], [16], [17], optimizing OSPF weights [32], and redundancy planning.

More recent studies [34], [35] target enterprise networks specifically. However, as described in Section I they employ the oversimplified “divide-and-conquer” strategy and perform network design in a stage-by-stage fashion. While these studies have advanced the state of the art of systematic network design, their models may produce designs with unnecessarily high configuration complexity, as we will elaborate in Section III. It is this unnecessary complexity that this work seeks to expose and minimize.

It should be noted that the recent progress in systematic network design owes largely to new abstractions from related work in several areas, including characterization of the designs of production networks (e.g., [23]), static analysis of network properties (e.g., [22], [38]), and the formulation of new configuration complexity metrics [7].

C. Software-Defined Networking

To combat network complexity, researchers have started investigating new software-defined networking (SDN) architectures based on logically centralized controllers and declarative configuration languages (e.g., Frenetic [15]). These approaches have the potential to simplify network design by

shifting complexity away from configuration of many individual devices to programming of few centralized controllers. However, we observe that SDN operators must carry out a similar design task of translating high-level reachability control requirements into flow rules. Since these flow rules will be installed on demand in the Ternary Content Addressable Memory (TCAM) of switches and once installed, checked for each packet passing through, it is desirable to minimize the number of such rules required. In this way, the design methodology and heuristics presented in this paper also apply to an SDN setting, as further discussed in Section VII.

III. AN INTEGRATED DESIGN FRAMEWORK

We now apply the integrated top-down approach as described in Section I to the user access control problem. With an illustrative example scenario, we identify the networking elements that are involved in realizing this important design objective, understand the source of configuration complexity of each element, and capture how the designs of individual elements may interact with each other and affect the overall complexity. We then present a framework for achieving an integrated design.

A. An Illustrative Example Scenario

Our example is based on the toy network shown in Figure 1a. There are two departments: Engineering and

Financial. Each department has users in multiple locations as shown. In addition, there is a set of servers. The access control policy is that the servers should only be accessed by Financial users. The following design steps are needed to implement the policy: (i) grouping the hosts into VLANs; (ii) assigning subnet addresses to VLANs; and (iii) installing a packet filter to restrict access to the servers. We are given the following design constraints: at most three VLANs can be created; and the available IP blocks are 10.0.1/24, 10.0.2/24 and 10.0.3/24.

Figure 1b illustrates a first possible design, where hosts are grouped into VLANs solely based on their *physical locations*. Unfortunately, this grouping scheme makes the packet filter configuration very complex: it is not possible to express the rules at the level of subnet prefixes for either VLAN 10 or VLAN 20, because they both contain hosts from two departments. As such, the filter rules have to be expressed at the level of individual IPs to permit Financial hosts in the two VLANs. This results in a large number of filter rules as illustrated in Figure 1b. We note that each packet filter is a sequential collection of filter rules, and each filter rule contains a pattern to be matched against packet headers, and an action (i.e., permit or deny) to be applied to packets whose header matches the corresponding pattern. The pattern part may be configured to match specific values of all or any subset of the five header fields: source and destination IPs and ports, and protocol; and thus creates static dependencies on those field values. Such dependencies must be manually configured and maintained and thus are a major source of configuration complexity.

Figure 1c depicts a different design, which ensures that hosts in the same VLAN belong to the same department. This design enables expressing filter rules at the level of subnet prefixes, and thus significantly reduces the number of rules. However, this design suffers from a different kind of configuration complexity: it requires configuring a large number of VLAN *trunk ports*, as denoted by the bold lines in the figure. Trunk ports are the switch ports that connect to other switches. Since each VLAN is a separate broadcast domain, it is important to properly constrain broadcast traffic to eliminate unnecessary broadcast overhead for increased performance and security. More specifically, every switch-to-switch link (called “trunk link”) must be configured to only allow traffic of appropriate VLANs. This is achieved by manually configuring the trunk ports to permit specific VLANs. Configuration of VLAN trunk ports is widely considered a major source of network configuration complexity, as operators must manually identify the correct set of VLANs to allow for each port. (Readers are referred to [34, Sec. II-A] for a more detailed explanation of VLAN trunk ports). A separate issue in this design is the way IP prefixes are assigned to VLANs: the prefix allocation scheme does not allow further aggregation of filter rules, as VLAN 10 and VLAN 30 (the two Financial VLANs to be permitted by the filter) are assigned non-aggregatable IP prefixes.

Figure 1d shows a third design. Same as the previous design, this design enables expressing filter rules at the subnet prefix level by ensuring that hosts in the same VLAN belong to the same department. However this design significantly reduces

the amount of VLAN trunk ports by grouping physically nearby hosts of the same department in the same VLAN. Furthermore, the IP prefix allocation scheme is also different: the two Financial VLANs are assigned aggregatable IP prefixes (10.0.1/24 and 10.0.2/24), which enables aggregation of filter rules, i.e., instead of using two rules “permit 10.0.1/24” and “permit 10.0.2/24,” we can use a single rule “permit 10.0.1/23” to permit both Financial VLANs.

We make two observations from this illustrative example. First, for the same target network, there exist multiple designs that are all *correct*. For our example network, there are at least 6 different designs (three different VLAN grouping schemes coupled with two different address allocation schemes.) However, different designs have different levels of configuration complexity. Second, the complexity of the resulting network is determined by both the VLAN configuration complexity (characterized by the number of trunk ports) and the packet filter complexity (characterized by the number of filter rules). Furthermore, the packet filter design is directly impacted by both the VLAN grouping scheme and the IP address allocation scheme. Thus, a design approach clearly will not work well if it treats VLAN grouping and IP allocation as total independent tasks and ignore their inherent interactions. For example, current top-down design approaches (e.g., [34], [35]) consider VLAN design as an isolated task, and thus they will solely seek to minimize the number of trunk ports without considering how doing so will impact the packet filter design, i.e., they will pick the first design shown in Figure 1b for our toy example, which is clearly not the best.

B. Interactions Between Designs of Multiple Elements

Realizing a host-level access control consists of designing the following networking elements: VLAN, IP address allocation, routing, and packet filters. We discuss below the role of each of those elements, their effect on configuration complexity, and the inherent interactions among their designs.

First, the VLAN design directly determines the number of trunk ports that need to be configured in the resulting network (and maintained during subsequent evolution). As explained above, VLAN trunk ports are a major source of operational complexity. Furthermore, the VLAN design also significantly affects the packet filter complexity, as it determines how hosts are grouped into subnets. Intuitively, if the VLANs align well with reachability policy boundaries (e.g., hosts in the same VLAN are subject to the same policy), then policy may be efficiently expressed at the level of subnet prefixes, resulting in a smaller number of filter rules. On the other hand, if VLANs are ill-aligned with policy boundaries (e.g., a single VLAN contains hosts subject to very different policy requirements), then filter rules may have to be expressed at the level of individual IP addresses, resulting in a large number of rules.

Second, the IP allocation scheme determines how filter rules may be aggregated and thus affects the number of filter rules in the resulting network. Intuitively, a good IP allocation scheme should minimize the number of filter rules by assigning aggregatable IP prefixes to VLANs that are subject to similar

access-control policy, so that a single filter rule can cover multiple VLANs by using an aggregated prefix. In contrast, IP allocation schemes that assign prefixes randomly or solely based on the physical location of VLANs is not likely to minimize filter rules. As another example, we have observed that some operational networks employ an IP allocation scheme that matches the third octet of every subnet prefix address to the corresponding VLAN ID, e.g., VLAN 100 will be assigned prefix $x.y.100.0/24$. These naive approaches treat IP allocation as an isolated design task and try to simplify the allocation scheme itself, but they fail to systematically consider how the IP allocation will affect the aggregation of filter rules.

Third, the routing design also affects the configuration complexity of packet filters. As a principle to ensure design correctness, if traffic between two subnets S_i and S_j is subject to filtering, then a filter must be placed on *every* possible layer-three path between the two subnets [35]. Routing design determines the layer-three topology, and thus directly affects the number of packet filters needed, the content of filters, and where they should be installed.

Scope of This Work: We observe that in practice, packet filters are typically only placed at the network edge, i.e., on the gateway routers of subnets. This design pattern has two major benefits. First, it guarantees that traffic to/from a subnet will always be filtered while simplifying the filter placement. That is, it relieves operators from having to find out all the possible layer-three paths between subnets. Second, all policies regarding a particular subnet can be implemented in a single location (i.e., its gateway router), which simplifies filter configuration and management. Based on this observation, in this paper we assume that packet filters will only be placed on subnet gateway routers. Given this assumption we no longer need to consider routing design, since the filter placement is now fixed and not affected by the layer-three topology. We do wish to acknowledge that systematic routing design is a challenging research problem on its own. We leave a more comprehensive investigation of designs where filters may be placed anywhere in the network to future work, and focus on VLAN design, IP allocation, and packet filter design in this paper.

C. Formulating the Access Control Design Problem

Following the integrated design methodology described in Section I, we now present a design framework for reachability-based access control, which integrates the design of the individual network elements identified above. In doing so, our goal is to enable the design process to be fully automated, while requiring only high-level specifications from operators. We first present a new abstraction that facilitates specifying and modeling reachability policy, and then present the framework.

1) *New Abstraction for Specifying Reachability Policy:* An essential input to our framework is the reachability control policy, and it is important to consider how it should be specified. The current “divide-and-conquer” design approach requires reachability policy to be specified at the VLAN/subnet level [35], i.e., it requires operators to specify a reachability matrix where each cell (i, j) denotes the reachability from

VLAN i to VLAN j . This abstraction works for the “divide-and-conquer” approach which assumes that the VLAN design has already been completed before designing packet filters. However, it does not work for our framework which it integrates the design of VLANs and packet filters, as VLANs themselves are to be determined by the solution. In addition, we believe that the VLAN-level reachability matrix is too low level as a policy abstraction, and it is tedious for operators to specify reachability policy using it.

In this work, we introduce a new abstraction for specifying reachability policy: a reachability matrix at “user role” level. We define a *user role* as a logical category that a set of users or servers belong to. Example user roles include faculty users, Computer Science users, financial servers, etc. Note that a user may have multiple roles, e.g., a CS professor can have both roles of CS users and faculty users. Each cell (i, j) of the reachability matrix specifies reachability policy from the user role i to the user role j . The advantage of this abstraction is that it allows policy to be specified at a higher level and independent of design.

2) *Design Formulation:* We formulate the design problem of reachability control as follows. We assume we are given the physical topology of the network, and the set of users/servers and their network locations. For each user/server, we are given its user roles. We are given the user-role-level reachability matrix as described above. Furthermore, we are given the maximal number of VLANs that can be created (denoted by N), and the available IP blocks. The design framework includes tasks of (i) mapping the set of users to at most N VLANs, (ii) assigning prefixes from the available IP space to the VLANs, and (iii) configuring packet filters to enforce the reachability policy. Our goal is to minimize the total configuration complexity of the resulting network. As discussed above, the configuration complexity (denoted as C_{total}) consists of VLAN-related complexity (denoted by C_v), measured by the number of trunk ports, and filter-related complexity (denoted by C_f), measured by the number of filter rules. Formally, we model the total configuration complexity as:

$$C_{total} = W_v * C_v + W_f * C_f \quad (1)$$

where W_v and W_f are the weight factors given to the two complexity categories, and can be customized by operators. For example, if operators of a network consider VLAN trunk ports more difficult to configure and maintain than filter rules, they can make W_v larger than W_f .

We observe that, while the VLAN grouping scheme and the IP allocation scheme both affect the configuration complexity of packet filters, VLAN design and IP allocation scheme are independent of each other, i.e., the design choices made in VLAN grouping won’t affect the available design space of the IP allocation scheme, and vice versa. Given this insight, we are able to formulate the design of reachability control as two joint design problems in order to make it more tractable:

- Joint design of VLANs and packet filters;
- Joint design of IP allocation scheme and packet filters.

We choose to perform the joint VLAN and filter design first, as the IP allocation design requires knowing the VLAN membership, i.e., which hosts belong to which VLANs,

in order to maximize prefix aggregation, as we will shown in Section V. The output of this joint design includes the VLAN grouping scheme, and an intermediate representation of packet filter rules expressed in terms of individual VLANs and hosts. This intermediate representation of packet filters then becomes part of the input to the second joint design. The output of the second joint design includes the IP allocation scheme, and final packet filters expressed in prefixes and individual IPs. In the next two sections, we formulate and solve the two joint design problems.

IV. JOINT DESIGN OF VLANs AND FILTERS

We first present models for formulating the joint design problem, and then develop heuristics for solving the formulated problem.

A. Formulating the Joint Design Problem

This design task is to map hosts to a set of VLANs and to derive packet filter rules expressed in terms of individual VLANs and hosts. There are several important considerations in doing so, as we detail below.

VLAN Count: The total number of VLANs that can be created in the design is determined by the hardware used in the network. This is because each VLAN runs its own instance of spanning tree, which consumes the memory and CPU resource of the switches. For example, a Cisco Catalyst 2950 switch can only support up to 64 spanning tree instances [11]. To model this constraint, we simply assume that operators will specify the maximum number of VLANs that can be created, which is denoted by N .

VLAN Size: A VLAN becomes a separate subnet at layer three, and thus the number of hosts that a VLAN can have is bounded by the size of the IP address block assigned to the corresponding subnet (assuming NAT is not used). For example, it is a common practice to limit the maximum size of a VLAN to that of a /24 subnet, i.e., at most 254 hosts. We assume the operators will specify the maximum VLAN size, denoted by MAX_VLAN_SIZE .

Correctness Criteria: To ensure the correctness of the design, the following two conditions must be satisfied. First, the given reachability policies must be correctly implemented through packet filters. Second, all hosts in the same VLAN must have full reachability toward each other, since they are all in the same broadcast domain.

Configuration Complexity: This design will determine the VLAN configuration complexity C_v (i.e., the total number of VLAN trunk ports in the resulting network). Further, it will also impact the packet filter configuration complexity. Note that the filter rules generated by this design are expressed in terms of individual VLANs and hosts. The VLANs and hosts will be assigned IP addresses in the second joint design, and thus the filter rules could be further aggregated when converted to the IP representation in that design. Thus, we model the total configuration complexity introduced by this design (denoted by C'_{total}) as follows:

$$C'_{total} = W_v * C_v + W_f * C'_f \quad (2)$$

W_v , W_f and C_v have been defined for Equation (1). C'_f is the configuration complexity of the packet filters generated by this design task, measured as the total number of filter rules. Clearly $C'_f \geq C_f$ as the joint design of IP allocation and packet filters may further reduce the number of filter rules through prefix aggregation.

Now we can formulate this joint design problem as follows:

Minimize: C'_{total}

Subject to:

- the correctness criteria, and
- the constraints on VLAN number and size.

B. Heuristics for Solving the Joint Design Problem

We present the details of our heuristics that work in a step-by-step manner. For ease of understanding, we use a running example to illustrate the algorithmic operations. The example network setup is shown in Figure 2a. There are eight user roles: Biology, Computer Science, IT, Faculty, Students, managers, operators, and servers. The reachability policy is also shown in the graph. We are given that $N = 6$, and $MAX_VLAN_SIZE = 254$, and $W_v = W_f = 1$.

1) Step 1 (Map Policy Groups to VLANs): As illustrated in Section III-A, it is often desirable for a VLAN to contain hosts subject to the same reachability policy, because doing so enables filter rules to be written at the level of an entire VLAN. To capture this insight in the design process, we leverage the abstraction of *policy groups* introduced by recent works [7], [33] including our own for network modeling. A policy group abstracts the set of hosts that are (i) subject to the same reachability policy towards other hosts and (ii) have full reachability among themselves. Clearly the set of policy groups forms a partition of all hosts. It is easy to see that a policy group is an atomic unit in deriving filter rules, i.e., if a packet filter allows traffic from one host in a policy group, it must also allow traffic from all the other hosts of the same policy group. Thus, the use of policy groups in the design process simplifies the reasoning of reachability control by allowing us to reason about groups of hosts together instead of individual ones. We believe the set of policy groups can be straightforwardly derived from the inputs of user roles and the role-level reachability matrix, but omit the details due to lack of space.

As a reasonable starting point of the design, we initially let each policy group become a separate VLAN. We then derive the filter rules. As mentioned in Section III-B, we have assumed that packet filters can only be placed on the gateway routers of the VLANs to be protected. Thus the filter rules can be determined in a straightforward way: for each VLAN, the corresponding packet filter permits all other VLANs (i.e., policy groups) that can communicate with this VLAN, according to the reachability matrix. We assume an implicit deny in the end of a packet filter, following the vendor convention. Filters that simply permit all traffic are omitted.

Figure 2b illustrate the design after this step. Seven policy groups are identified straightforwardly from the inputs: CS faculty (shown as CS-F on graph) which resides in two different locations, CS students (CS-S), Biology

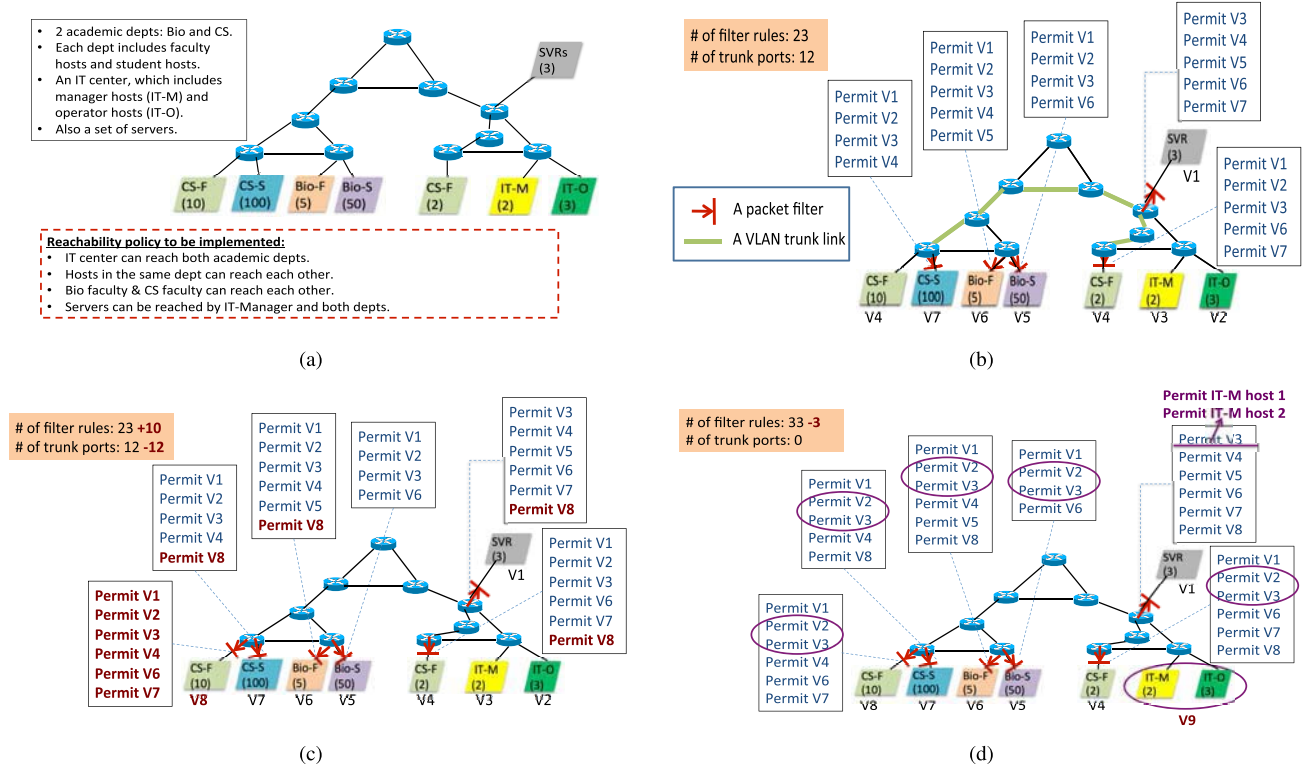


Fig. 2. An running example for illustrating the operations of our heuristics for joint design of VLAN and packet filters. (a) The example network and its reachability policies. (b) After step 1, each policy group becomes a VLAN. Packet filters and trunk ports are determined accordingly. (c) After step 2, the original $V4$ is partitioned into two VLANs: the new $V4$ and $V8$. This reduces C'_{total} by 2, by eliminating all the 12 trunk ports, though incurring 10 additional filter rules. (d) In step 4, the heuristics choose to combine the old $V2$ and $V3$ to form the new $V9$, as doing so reduces C'_{total} by 3.

faculty (Bio-F), Biology students (Bio-S), IT managers (IT-M), IT operators (IT-O) and servers (SVR). Each policy group has been placed in a separate VLAN. For example, the entire CS-Faculty policy group becomes VLAN $V4$. The corresponding VLAN trunk ports to be configured are shown by the bold links connecting those ports. The packet filters are also shown, and as expected all filter rules are expressed at the VLAN level. Finally, the amount of configuration complexity in terms of filter rules and trunk ports after this step is also shown.

2) *Step 2 (Selectively Partition VLANs With Large Span)*: For each VLAN created in Step 1, we now evaluate whether it is beneficial (i.e., leading to smaller C'_{total}) to partition it into two smaller VLANs. If so, we will execute the partitioning, and iteratively evaluate for the resulting two smaller VLANs. We repeat this step for every VLAN until we cannot further reduce C'_{total} by partitioning existing VLANs. Our insight for this step is as follow.

On one hand, partitioning a VLAN that has a large span could potentially reduce C'_{total} as it could significantly reduce the number of trunk ports (i.e., C_v). Consider $V4$ (CS-Faculty policy group) in Figure 2b as an example. By partitioning it into two smaller VLANs, i.e., the new $V4$ and $V8$ in Figure 2c, we eliminate the need for any trunk port for this VLAN, and thus reduce C_v . On the other hand, partitioning a VLAN could also potentially increase C'_{total} as it could lead to more filters and/or filter rules required, i.e., an increase in C'_f . There are two reasons for this. First, after the partitioning it may be necessary to install a *new* packet filter to protect a newly created VLAN. For example, in Figure 2c there is a new packet

filter that protects the newly created $V8$, which introduces 6 new rules. Second, it may be necessary to add additional rules in the *existing* filters, to permit a newly created VLAN. For example, in Figure 2c a rule “permit $V8$ ” is added to four existing filters.

More specifically, we employ the K-means clustering algorithm (with $K = 2$) to decide how a VLAN should be partitioned into two, such that the reduction in C_v is maximized. In configuring the clustering algorithm, we let each host in the VLAN be a node, and the distance between two nodes be the length of the shortest layer-two path between the corresponding hosts. The clustering algorithm then groups nearby hosts into the same VLAN and thus minimizes the need of trunk ports.

For our running example, we find that by partitioning the old VLAN $V4$ in Figure 2b into two smaller VLANs $V4$ and $V8$ in Fig. 2c, we reduces C_v (i.e., the number of trunk ports) by 12, but increases C'_f (i.e., the number of filter rules) by 10. As we assume $W_v = W_f = 1$, the total complexity is reduced by 2, according to Equation (2). Hence, we execute the partitioning since it is beneficial to do so. We also find that it is not beneficial to partition any other VLAN. Figure 2c shows the resulting design after this step.

3) *Step 3 (Partition VLANs With Too Many Hosts)*: This step ensures that the constraint on VLAN size is met. It checks each VLAN in the current design to see whether it contains more hosts than the specified MAX_VLAN_SIZE . If so, it again uses the K-means clustering algorithm described in the previous step to partition the VLAN into two. This process

iterates until all VLANs have been reduced to a size no larger than the MAX_VLAN_SIZE .

For our running example, Since none of the VLANs contains more than 254 hosts, this step will not partition any VLAN.

4) *Step 4 (Selectively Combine VLANs)*: This step has two purposes: further reducing the total complexity C'_{total} , and also ensuring that the constraint on the VLAN count is met. It achieves both by selectively combining pairs of VLANs in an iterative process as described below.

For every *eligible* pair of VLANs, the heuristics evaluate the complexity impact of combining them. A pair of VLANs is eligible to be combined if (i) the sum of the hosts in both VLANs is not greater than MAX_VLAN_SIZE , and (ii) the hosts in both VLANs have full reachability toward each other. For every eligible VLAN pair, we calculate the potential change in C'_{total} if the two were combined into a single new VLAN. We then select the pair with the maximum reduction in C'_{total} to execute the combining. We repeat this process until the following two conditions are *both* met:

- The total number of VLANs is not greater than N (i.e., the maximum number of VLANs that can be created); and,
- It is not possible to further reduce C'_{total} by combining any more eligible pair of VLANs.

To understand why combining VLANs could possibly lead to reduction in C'_{total} , consider $V2$ and $V3$ in Figure 2c as an example. If we combine those two VLANs into the new $V9$ as illustrated in Figure 2d, then for all packet filters that need to permit both $V2$ and $V3$ by using two separate rules, they now only need to permit the new $V9$ using a single rule, leading to a reduction of rules.

However, this benefit does not come without potential penalty. The penalty is two-fold. First, if there is any packet filter that permits only one of the two original VLANs, then it cannot permit the combined new VLAN. For example, in Figure 2c the packet filter protecting $V1$ (i.e., the servers) only permits $V3$ but not $V2$. So after $V2$ and $V3$ are combined to form the new $V9$ as shown in Figure 2d, the filter cannot simply change to permit $V9$ instead, because doing so would wrongfully grant access to hosts in the original $V2$. Hence, the filter now has to permit *individual hosts* in $V3$ as shown in Figure 2d, leading to an increase in the number of filter rules. Second, combining two VLANs could also require configuring additional VLAN trunk ports, if the two VLANs are in different locations. Though in our example this is not the case as $V2$ and $V3$ connect to the same switch.

For our running example, the heuristics will choose to first combine $V2$ and $V3$ to form the new $V9$, because doing so results in a reduction in C'_f by 3 while keeping C_v unchanged. This is illustrated in Figure 2d. In fact this is the only pair of VLANs that will result in a reduction in C'_{total} if combined. All other VLAN pairs when combined will cause C'_{total} to increase. However, since the total number of VLANs after combining $V2$ and $V3$ is 7, which is greater than the given limit of $N = 6$, another pair of VLANs has to be combined. The heuristics will again evaluate all eligible pairs and then choose to combine $V1$ and $V4$ to form the new

VLAN $V10$, as doing so results in the least increase in C'_{total} (C_v and C'_f will be increased by four and two respectively). After that, both conditions listed above are met and this step stops.

C. Complexity of Algorithm

The algorithmic complexity of Step 1 is $O(p^2)$ where p is the number of policy groups in the network, because it needs to derive filter rules for each pair of policy groups. Since the number of policy groups is bounded by the number of hosts n , the complexity may be considered as $O(n^2)$. The complexity of Steps 2 and 3 is dominated by the complexity of the k-means clustering algorithm. Although k-means clustering is an NP-hard problem, there exist efficient heuristics that run in approximately $O(bkdi)$ [37], where b is the number of observations (in our algorithm this is the number of hosts in the VLAN to be partitioned, which is bounded by n , the total number of hosts), k is the number of clusters (always set to 2 in our algorithm), d is the number of dimensions in measuring distance between observations (this is again the number of hosts in the VLAN to be partitioned and bounded by n), and i is the number of iterations needed until convergence, which is said to be often small, and results only improve slightly after the first dozen iterations [37]. Since Steps 2 and 3 run the k-means clustering algorithm for at most every VLAN, and clearly the number of VLANs is bounded by n , the overall complexity of these two steps may be considered as $O(n^3i)$. Finally, the complexity of Step 4 is $O(n^3)$ since it will at the most consider combining every pair of VLANs.

V. JOINT DESIGN OF IP ALLOCATION AND FILTERS

We first formulate the joint design problem, and then present a heuristic solution based on finding the maximum weighted matching on a graph. In describing the heuristics, we continue to use the same running example from the previous section.

A. Formulating the Joint Design Problem

This design task is performed after the joint design of VLANs and packet filters that is presented in the previous section. The inputs are: (i) the VLAN grouping scheme; (ii) the packet filters in the intermediate representation (i.e., expressed in terms of individual VLANs and hosts); and (iii) available IP blocks. The goal of this design is to find a good scheme of allocating IP prefixes to VLANs such that the resulting number of filter rules is minimized.

For example, in Figure 3a, the packet filter that protects $V6$ contains four rules to separately permit $V5$, $V8$, $V9$ and $V10$. However, if two aggregatable prefixes (say 10.0.1/24 and 10.0.2/24) are assigned to $V5$ and $V8$, then the two VLANs can be permitted together in one rule that permits the aggregated prefix 10.0.1/23. Even better, if prefixes 10.0.3/24 and 10.0.4/24 are also assigned to $V9$ and $V10$, then further aggregation can be achieved and the filter will need only a single rule “permit 10.0.1/22” to permit all four VLANs. Further, since multiple packet filters are typically involved, the address allocation scheme should prioritize the assignment of aggregatable prefixes based on how frequently the candidate VLANs appear together and receive the same treatment

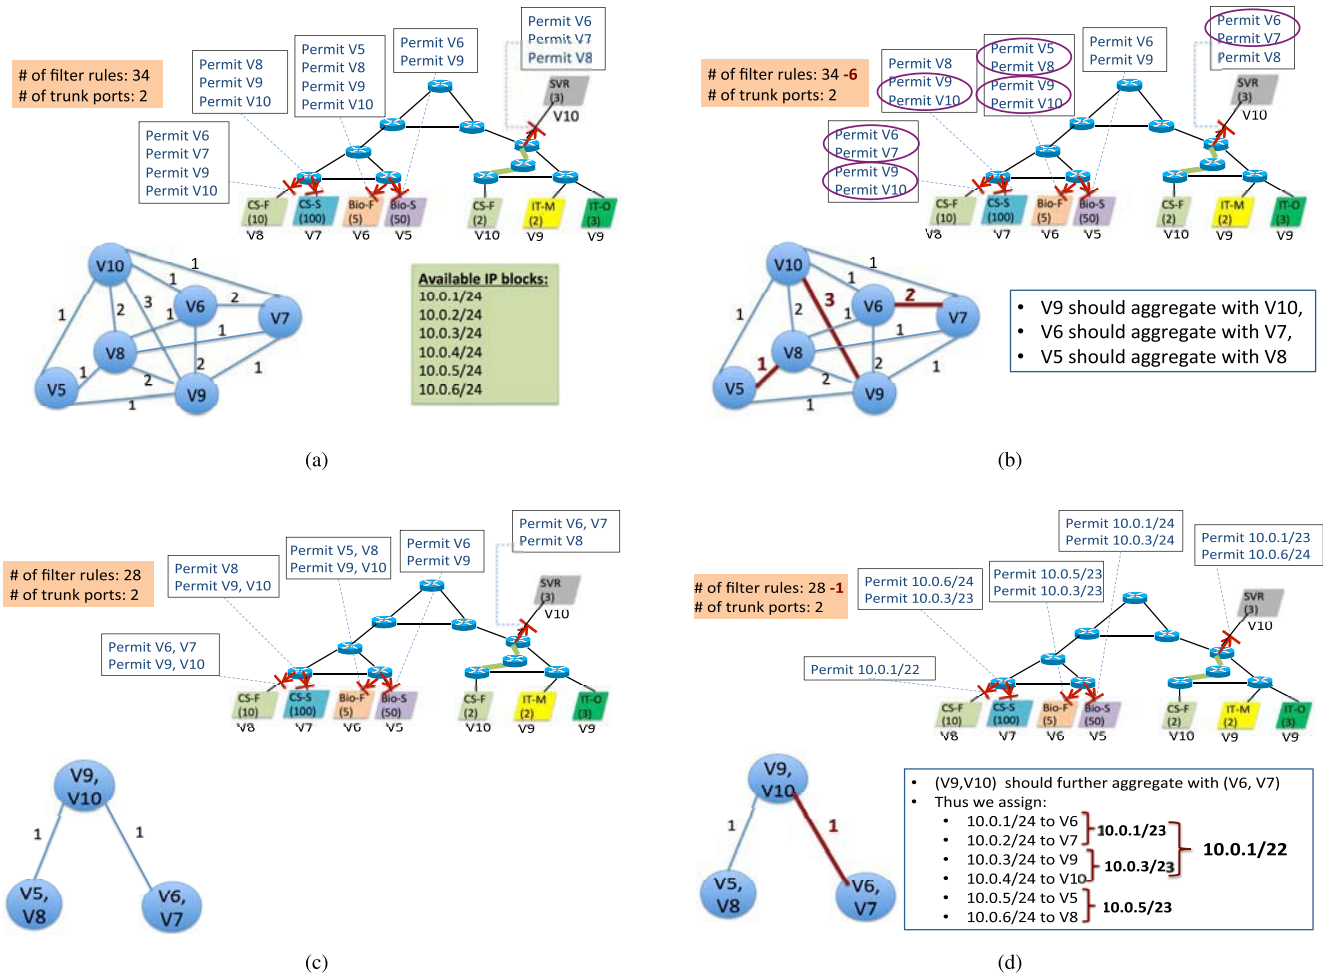


Fig. 3. The example continued from previous section for illustrating operations of our heuristics for assigning IP prefixes to VLANs. (a) Construct G_{24} for the example network, after the VLAN design. (b) The maximum weighted matching of G_{24} . (c) Construct G_{23} based on the results in the previous iteration. (d) The final prefix allocation scheme based on the maximum weighted matching on both G_{23} and G_{24} .

(i.e., permit or deny) in all filters, in order to minimize the total number of filter rules.

Our focus in designing the address allocation scheme is on the VLAN/prefix level, i.e., we focus on assigning IP prefixes to VLANs. We do not consider how IPs should be assigned to *individual* hosts *inside* each VLAN. Even though carefully assigning IPs to individual hosts could potentially enable intra-VLAN IP aggregation and reduce the number of filter rules concerning individual hosts, we believe this may be too low level that it is impractical to require operators to configure and track individual IP assignment. In practice, DHCP is often used so that hosts will receive their IPs automatically from the IP blocks assigned to their VLANs/subnets. Thus, we simply assume that individual IP assignment to hosts is done randomly, and do not consider possible aggregation of individual host IPs. We do wish to note that the heuristics presented below can be directly applied to finding the individual host IP assignment, if the operators wish to do so.

Correctness Criteria: To ensure the correctness of the design, the following two conditions must be satisfied. First, the prefixes assigned to VLANs must be chosen from the pool of available IP blocks. Second, when aggregating filter

rules, the given reachability control policies must be correctly implemented.

Configuration Complexity: This design will determine the final packet filter complexity C_f (i.e., the total number of filter rules in the resulting network).

Formally, the joint design of prefix allocation and packet filters can be formulated as follows:

Minimize: C_f

Subject to: the correctness criteria.

B. Heuristics for Solving the Joint Design Problem

The key idea for solving the problem is to model it as *finding the maximum weighted matching on a graph*. Our solution works in iterations over the prefix lengths, starting from /32, then /31, then /30, and so on (i.e., in each iteration the prefix length is decreased by 1). In the iteration that concerns prefix length l , we construct a graph G_l whose vertices are prefixes of length l that are available and can be assigned to VLANs. There is an edge between two vertices, if the corresponding VLANs *appear together and receive the same treatment* in at least one filter. The weight of an edge is defined as the

number of filters in which the two corresponding VLANs appear together and receive the same treatment.

To illustrate, Figure 3a shows our running example continued from the previous section. Note that only rules concerning an entire VLAN are shown here, and rules concerning individual hosts are omitted, as the focus here is on assigning prefixes to VLANs. The graph G_l is empty for iterations concerning prefix length from $/32$ to $/25$, as the VLANs in this particular example are all of $/24$. In the iteration concerning $/24$, the corresponding graph G_{24} is shown in the lower left part of Figure 3a. The numbers shown on the edges are the weights. For example, the weight of the edge $(V9, V10)$ is 3, as the two VLANs being permitted together in three packet filters.

Now the design problem of finding the best allocation scheme of prefixes of length l can be solved by finding the maximum weighted matching on G_l . A *matching* on a graph is defined as a subset of edges such that none of them share a common vertex. The *maximum weighted matching* is defined as a matching for which the sum of the weights of the matched edges is as large as possible. We note there exist efficient algorithms (e.g., [5]) that take polynomial time to find the maximum weighted matching on a general undirected graph. Now in our context, for each edge included in the maximum weighted matching, the corresponding two VLANs should be assigned aggregatable prefixes. It is easy to see that by maximizing the weight of the selected matching on G_l , we maximize the opportunity to reduce the number of filter rules through prefix aggregation.

We leverage the algorithm described in [5] to find the maximum weighted matching for our running example, and the result is marked in red on the G_{24} graph shown in Figure 3b. According to the result, we should assign aggregatable prefixes to $V9$ and $V10$, to $V6$ and $V7$, and to $V5$ and $V8$. Doing so will reduce the number of filter rules by 6 as illustrated in Figure 3b.

The process of constructing G_l and finding the maximum weighted matching on it continues for larger prefixes. It stops when G_l does not have any edge. For our example, after the above iteration of $/24$, we are left with three $/23$ prefixes, which are aggregated prefixes of the corresponding pairs of VLANs as specified by the maximum weighted matching in the previous iteration of $/24$. So the new graph G_{23} can be constructed as shown in Figure 3c. G_{23} has three vertices, which are the three VLAN pairs each receiving aggregatable prefixes in the previous iteration. There is an edge between two vertices, if all four involved VLANs appear together and receive the same treatment in at least one filter. The weight of an edge is the number of filters in which all four involved VLANs appear together and receive the same treatment. For the running example, either edge could be the maximum weighted matching for G_{23} , and our heuristics will randomly pick one, say the edge $(\{V9, V10\}, \{V6, V7\})$ as shown in Figure 3d. This means that aggregatable $/23$ prefixes will be assigned to the two pairs of VLANs $\{V9, V10\}$ and $\{V6, V7\}$, so that they can be further aggregated to a $/22$ prefix. As a result, whenever a filter needs to permit all those four VLANs, it can simply permit the aggregated $/22$ prefix in a single rule. The process stops after $/23$ for the example network.

Based on these results, for our running example it is best to assign $10.0.1/24$ to $V9$, $10.0.2/24$ to $V10$, $10.0.3/24$ to $V6$, $10.0.4/24$ to $V7$, $10.0.5/24$ to $V5$, and $10.0.6/24$ to $V8$. This prefix allocation scheme reduces the number of filter rules by 7, which is the maximum reduction that can be achieved through prefix aggregation. The final design is shown in Figure 3d.

C. Complexity of Algorithm

It is easy to see that the complexity of this algorithm is dominated by the complexity of finding the maximum weighted matching, which is $O(v^3)$, where v is the number of vertices in the graph. In our context, each vertex corresponds to a VLAN, whose number is bounded by the total number of hosts n . The maximum weighted matching algorithm will at most be executed 31 times (from prefix length $/32$ to $/1$). Hence the overall complexity of this algorithm is $O(n^3)$.

VI. EVALUATION

We evaluate our integrated design framework using two medium-sized university campus networks (termed “*university-1*” and “*university-2*” throughout this section). Each network is assigned a $/16$ IP space. For *university-1*, our dataset includes configuration files of all devices, as well as the complete layer-2 topology data obtained through Cisco Discovery Protocol (CDP). For *university-2*, our dataset includes configuration files of all devices, however the layer-2 topology information is not available to us; as such it cannot be used to evaluate our VLAN design. Thus we will use the *university-1* dataset to evaluate the complete design framework, and use the *university-2* dataset for evaluating the address allocation only.

VLANs are extensively used in both networks: *university-1* uses a total of 69 VLANs, with the vast majority of them assigned a $/24$ prefix address; *university-2* uses a total of 348 VLANs with smaller size – most are assigned a $/26$ prefix. A large number of packet filters (i.e., access-control-list, or ACLs) is present in both networks as well. The vast majority of filters in both networks are installed on the gateway routers of VLANs only, and there is no filters in the network core. This matches our assumption of filter placement (see Section III-B) very well.

A. Characterizing Policy Groups

Although our design framework allows operators to specify reachability policies using the user-role-level reachability matrix (Section III-C1), unfortunately for both campus networks under study a complete and up-to-date documentation of all reachability policies is not available. Thus, we take an alternative approach and reverse engineer all the policy groups based on the device configuration files, using the methodology presented in [8]. We are able to identify all the policy groups in both networks. *University-1* and *university-2* contains 116 and 58 policy groups respectively, and furthermore, make the following interesting observations about them.

First, for both networks, the majority of policy groups are very small with only a couple of hosts; however, there are

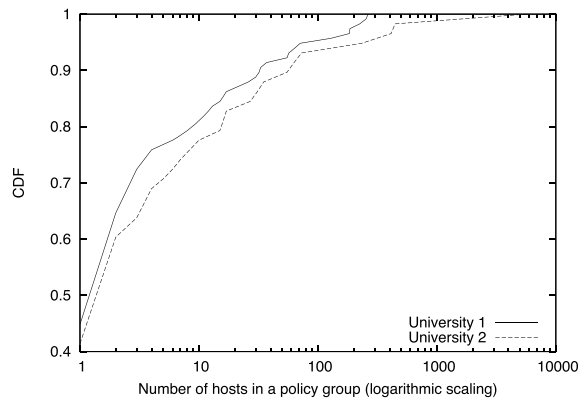


Fig. 4. CDF on the number of hosts in each policy group.

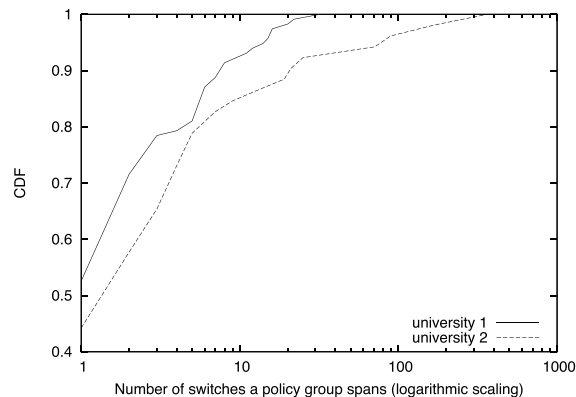


Fig. 5. CDF on the number of switches that a policy group spans.

some policy groups that are quite large. The size distribution of the policy groups are shown in Figure 4. In *university-1*, the largest policy group includes 264 hosts which is larger than any single VLAN. 23 policy groups contain 10 hosts or more, and 10 of them contains 50 hosts or more. Further investigation shows that many of the small policy groups are servers, special purpose (e.g., VoIP) boxes, or management hosts (e.g., operators granting their own office desktops special privilege so that they can log on to remote switches and routers right from their office.) On the other hand, the largest policy groups are student dorm hosts and classroom PCs. These hosts of large volume and span many buildings, but are all subject to the same reachability policies. In *university-2*, the largest policy group contains the vast majority (5690) hosts. 15 policy groups contain 10 hosts or more, and 7 policy groups contains 50 hosts or more.

Second we investigate the footprint of these policy groups, by measuring the number of switches they span. More specifically, for each policy group we measure the number of switches that one or more of its hosts directly connect to. The results are summarized in Figure 5. While around half of the policy groups in both networks connect to only a single switch, 20% (27%) of them span 5 or more switches, and 10% (15%) span 10 or more switches, for *university-1* (*university-2*). The largest policy group in *university-1* spans 32 switches, which turns out to be the dorm machines. The largest policy group in *university-2* spans 362 switches. Next, We measure for every switch the number of policy groups that connect

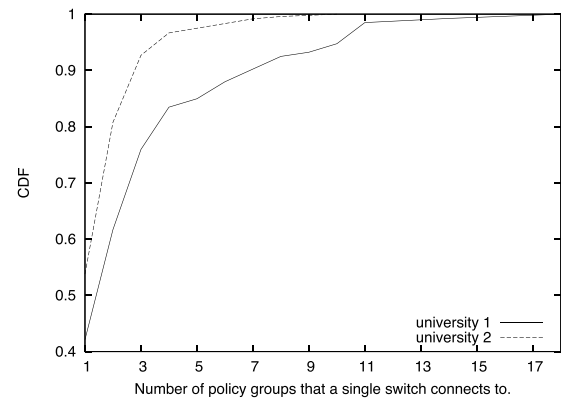


Fig. 6. CDF on the number of policy groups that a switch connects to.

to the switch. The results are summarized in Figure 6. For *university-1*, while 42% of the switches connect to entirely one policy group, 24% of the switches connect to 4 or more policy groups, and 12% of the switches connect to 7 or more policy groups. The maximum number of policy groups that a switch connects to is 18. For *university-2*, 54% of the switches connect to entirely one policy group, 7% connect to 4 or more policy groups, and 2% connect to 7 or more policy groups. The maximum number of policy group that a single switch connects to is 10.

These results show that the “divide-and-conquer” approach [34], [35] that solely seeks to minimize the number of VLAN trunk ports will not work well for minimizing the overall configuration complexity. This is because that approach will group hosts purely based on their physical locations (i.e., the switches they connect to) and thus will likely place multiple policy groups connecting to the same switch in the same VLAN. This is particularly true when the number of VLANs that can be created is smaller than the number of policy groups, which is the case here. As a result, many filter rules will have to be expressed at the individual IP level, resulting in a large number of rules as illustrated by the example design in Figure 1b.

B. Evaluating the Joint Design of VLANs & Filters

We now evaluate the effectiveness of our heuristics for a joint design of VLANs and packet filters. We use the *university-1* dataset throughout this study. (we could not use *university-2* dataset as we do not have the layer-2 topology data for that network).

In the *university-1* campus network, packet filters are placed on over 70 layer-3 switches and routers, with more than 6000 rules in total. It is surprising to find that the majority of those rules are at the individual IP level. On the other hand, the current network also has a large number (2500+) of VLAN trunk ports, needed to connect hosts in different physical locations into the same VLAN. Due to these facts, the current network has a high degree of configuration complexity. We do wish to note that the campus network is well managed by a dedicated team of highly skilled operators, and that many hours of design time have been spent on finding a solution to reachability control. We believe that these observations confirm that for a large-scale network with fine-grain reachability




Fig. 7. Comparing the number of filter rules and the number of VLAN trunk ports produced by our joint VLAN/filter design heuristic to those by the ad-hoc design approach and by a divide-and-conquer design.

control requirements, it is just too difficult for any operators to manually search for the design that minimizes the complexity. This highlights the need for our integrated top-down design approach.

To execute our design heuristics, we set N (the maximum number of VLANs allowed in the design) to be 69, the same number of VLANs used in the current network. The MAX_VLAN_SIZE is set to 254, which also matches the current network design. We set the weight factors W_f and W_v (for packet filter complexity and VLAN trunk port complexity, respectively) to be 1. We then run our heuristics (implemented by a set of Perl scripts) on our dataset. The heuristics run sufficiently fast and complete the design in less than two minutes on a PC with a quad-core i7 CPU. The resulting configuration complexity is shown in Figure 7 and discussed below.

Comparison With the Ad-Hoc Design in Operation:

The first and third clusters of bars in Figure 7 correspond to the ad-hoc design approach that operators used to produce the current network, and our joint VLAN and filter design approach, respectively. In each cluster, the two bar shows the total number of filter rules and VLAN trunk ports resulted from the design. The results show that (i) our framework effectively reduces the total number of filter rules down to 1809, which is only 30% of the number of filter rules in the current network; and (ii) our framework also reduces the number of VLAN trunk ports down to 716, which is only 29% of the number of trunk ports in the current network.

Comparison With the State-of-the-Art Divide-and-Conquer Approach: We applied to the same network the “divide-and-conquer” design approach that we previously developed [34], [35]. We believe that this is the state of the art in VLAN design, as it has been widely accepted by the research community and adopted by many follow-up works (e.g., [18], [24], [25]). In principle, this approach solely minimizes the span of VLANs, and thus minimizes the number of trunk ports, by grouping physically nearby hosts together (as long as the hosts can reach each other, which ensures the correctness criterion is met). Hence the approach would correspond to the design method illustrated by Figure 1b. As the approach

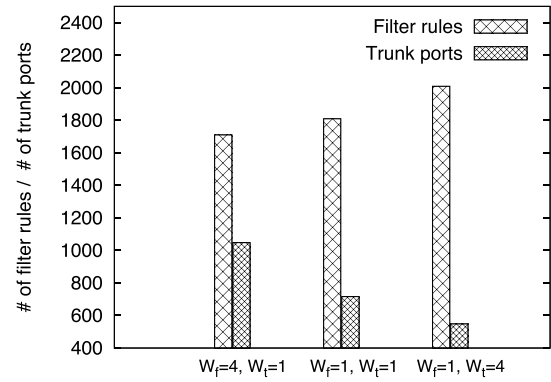


Fig. 8. Sensitivity of the results to the W_f and W_v parameters.

does not consider how host grouping may affect packet filter design, and based on the observations from Section VI-A, it will likely place hosts from different policy groups in the same VLAN, resulting in a large number of fine-grained filter rules as illustrated by Figure 1b. This insight has been confirmed by the evaluation results, represented by the second cluster of bars in Figure 7. While this “divide-and-conquer” approach achieves the least number of VLAN trunk ports among all the design approaches, it creates by far the most filter rules, nearly tripling the number of the ad-hoc design. In comparison, the number of filter rules created by our integrated design approach is only 10% of the “divide-and-conquer” approach, a reduction of over 16000 rules, at the modest price of having a couple hundred more trunk ports.

Overall, the above results clearly show that our joint VLAN and filter design achieves significantly better host grouping, and results in substantially lower configuration complexity, compared to both the ad-hoc design used by operators and the start-of-the-art “divide-and-conquer” design.

We next study the sensitivity of the results to the W_f and W_v values. For this purpose, we consider two alternative design scenarios. In the first scenario, the complexity of configuring VLAN trunk ports is considered four times higher than that of configuring filter rules, and thus we set $W_f = 1$ and $W_v = 4$. In the second scenario, the complexity of configuring filter rules is considered four times higher than that of configuring VLAN trunk ports, and thus we set $W_f = 4$ and $W_v = 1$. We run the heuristics with these two additional setups on the same dataset, and Figure 8 summarizes the results. Each cluster of bars corresponds to a specific choice of W_f and W_v . In each cluster, the first bar shows the total number of filter rules in the resulting design, and the second bar shows the total number of VLAN trunk ports. We make two observations. First, For all settings, the total configuration complexity is substantially lower than that of the current network. This shows that our heuristic effectively reduces the complexity regardless of the choice of W_f or W_v values, and thus can be applied to a wide range of design scenarios. Second, the results also show that our heuristic can intelligently trade off the two complexity factors for different design scenarios, and produce the best design for




Fig. 9. The number of filter rules produced by the current ad-hoc design, the joint design of VLANs and filters alone, and the full integrated design framework including the prefix allocation step.

each scenario. For example, when VLAN trunk ports are given a higher complexity weight, the produced design uses less trunk ports, at the cost of using more filter rules. In contrast, when packet filters are given a higher complexity weight, the produced design uses fewer filter rules, at the cost of more VLAN trunk ports.

C. Evaluating the Joint Design of IP Allocation & Filters

We next evaluate the effectiveness of our heuristic for a joint design of IP address allocation and packet filters. Recall that our heuristic takes as input the packet filters produced in the joint design of VLANs and filters, and those filters are expressed in terms of individual VLANs and hosts. Note that the focus of this paper is on allocating prefix addresses to VLANs, and the heuristic only designs the IP allocation scheme at the prefix level (Section V-A). For this purpose, we continue to use the *university-1* dataset. Again our heuristic runs relatively fast and completes the design in less than one minute. The result is shown in Figure 9. This figure shows the number of filter rules in three designs: (i) the current network, (ii) the joint design of VLANs and packet filters alone, with $W_f = W_v = 1$, and (iii) the full integrated design including the IP prefix allocation step. (We do not show the VLAN trunk port data in this figure since they are not impacted by the prefix allocation scheme.) We see that by integrating the prefix allocation design and the packet filter design, our heuristic is able to further reduce the total number of filter rules down to 841. This halves the total number of filter rules (including both VLAN-level and host-level rules) produced by the joint design of VLANs and packet filters alone, and is only 14% of the number of filter rules in the current network. Together, the total amount of configuration complexity (including both filter rules and VLAN trunk ports) incurred by our integrated top-down design approach is only 18% of that incurred by the current ad-hoc design approach. Overall, these results demonstrate the effectiveness of our integrated design approach in reducing network configuration complexity.

We also consider a separate evaluation of the joint IP allocation and filter design, that focuses entirely on this part of the design framework alone. We keep the current VLAN grouping of both campus networks unchanged, and only designs the




Fig. 10. Filter rule reduction resulted from designing the IP allocation alone (without touching the VLAN grouping scheme), for both campus networks.

IP prefix addressing scheme, using the algorithm presented in Section V. In doing so, we honor the size of all the subnet addresses in the current assignment, i.e., if currently VLAN x is given a $/24$ subnet address p , we may change the address from p to q from the IP space of the campus network, but the size of q must also be $/24$. Figure 10 summarized the results. For *university-1*, by considering the IP allocation alone (without touching the VLAN grouping), our design approach is able to reduce the number of filter rules by 29% by enabling better aggregation. For *university-2*, our approach resulted in a 34% reduction in filter rules. Overall, these results further confirm the effectiveness of the joint IP allocation and filter design.

VII. DISCUSSION AND OPEN ISSUES

Applying the Integrated Approach to Other Operation Objectives: We consider the presented design methodology as briefly outlined in Section I an important contribution to the network design problem in its own right. Even though the problem formulations, algorithms and heuristics developed in this paper are specific to the design of access control, the integrated design methodology is general, neither limited to nor tied to that specific context. In fact, we observe that virtually every network operation objective involves designing multiple networking elements; therefore, we expect our integrated methodology to have wide applicability in top-down network design.

Take QoS design for instance. A QoS solution typically involves end-to-end traffic engineering (e.g., through routing) and per-link traffic management (e.g., through policing, marking and shaping of packets on routers) [36]. Intuitively, a more sophisticated routing design such as routing traffic of different QoS classes over different paths, which has its own complexity to implement and maintain, can achieve a more predictable and simpler traffic pattern at each router, and subsequently simplify the per-link traffic management. On the other hand, a simplistic routing design that allows all classes of traffic on all links would likely complicate the traffic management task. Thus, when designing a QoS solution, it is important to jointly consider routing and traffic management in order to minimize the overall complexity.

Considering Other Aspects of Network Complexity: Our integrated design methodology is not tied to the configuration complexity metric used in this paper. In principle,

the approach will work with any complexity metric that is quantifiable based on design parameters. For example, Chun *et al.* [10] have proposed to measure the amount of dependencies between states maintained at different routers. Conceivably, the collection of states required at each networking device can be inferred from the choices of protocols and other such decisions at design time. If that is indeed the case, one may use the state-centric metric in place of the configuration-centric metric in formulating new optimization problems similar to those presented in Sections IV and V. An interesting open question is how much the set of optimal design choices would vary from metric to metric.

It is noteworthy that our literature search is unable to identify additional complexity metrics subject to the criteria of being (i) objectively quantifiable and (ii) directly linked to design choices. On a positive note, the networking community is increasingly aware of the importance of developing formal models and metrics for defining and quantifying network complexity. A new group has been formed within the Internet Research Task Force (IRTF) to specifically promote research in this direction. Particularly of interest is the call by this group to develop high-level complexity metrics to help design networks with more predictable behaviors and less resembling of complex nonlinear systems where a small local perturbation may lead to a cascading system wide failure [6].

Applying the Integrated Design Approach to Evolve Existing Networks: While the focus of this paper is on designing new (i.e., “green-field”) networks, we believe that the integrated approach presented here can be extended to evolve existing “brown-field” networks as well. Our prior work [34] on network evolution shows that, when making changes to their networks, operators typically face many design choices. For example, in that work [34] we have shown that the majority of changes to a network is adding new hosts, and that for this kind of changes, operators have the choice of either adding them to one or more existing VLANs, or creating one or more new VLANs; in the latter case, operators must also make the design decision of what IP prefixes should be assigned to the newly-created VLANs. Similarly, moving existing hosts to new locations is another common type of changes, and operators need to decide for the hosts being moved whether it is beneficial operationally to change their VLAN membership. We note that the design decisions required for these common change events are quite similar to the design decisions studied in this work; the main difference is that, for making changes to an existing network, the existing design (such as existing VLAN design and IP allocation scheme) must be modeled as additional constraints to the formulated optimization problem. We are thus confident that the integrated design framework presented in this work can be easily extended to evolve existing networks. A complete investigation is out of the scope of this paper and is the subject of our future work.

Applying the Integrated Approach to Optimize SDN Flow Rule Generation: Recent research [20], [27], [29] on SDN advocates that the controller platform should provide a “one big switch” abstraction to the applications running on it. This abstraction enables the application programmers to specify policies at a high level (i.e., network level) and let the

controller compile those policies into low-level (i.e., switch-level) flow rules and install them on individual switches. In doing so, a fundamental constraint is the limited TCAM space on the commodity switches where the rules will be stored. Thus, it is desirable to minimize the number of flow rules. Existing proposals on this front again take a simplified “divide-and-conquer” approach and assume that the IP allocation scheme has been decided before generating and distributing the rules, even though how the IP addresses are assigned can significantly affect how flow rules may be aggregated. We believe that our integrated methodology can be applied to jointly design IP allocation and rule generation and distribution to minimize the resulting number of rules. The heuristics presented in this paper may be leveraged in that context as well. We leave a thorough investigation in this direction to future work.

Optimality vs. Tractability: We consider the formulations and algorithms presented in this paper only one candidate solution of a spectrum of possible integrated design frameworks for reachability-based access control. For example, it may be feasible to formulate VLAN design, IP address allocation and routing design into a single optimization problem. Broadly speaking, we observe that two competing factors, *optimality* (in terms of how many network elements are unified) and *tractability* (whether a practical solution can be found), are at play with the integrated approach. An interesting open question is whether a class of design points (“sweet spots”) exists that strikes the right balance between the two factors.

VIII. CONCLUSIONS AND FUTURE WORK

We have shown the importance and effectiveness of an integrated top-down network design methodology for systematically identifying, among all designs that meet a given operational objective, the one(s) with the minimum configuration complexity. The approach enables us to rigorously formulate two new optimization problems as part of a design framework for accomplishing a network’s access control policy while avoiding unnecessary configuration complexity. The power of the new formulations comes from a unified model that captures the intricate interplays between design decisions concerning VLANs, IP addresses, and packet filters. While this paper focused on reachability-based access control as a concrete application, we believe that the integrated design methodology is applicable not only to a variety of design objectives and tasks for today’s networks, but also to the emerging SDN paradigm.


To the best of our knowledge, this is the first work that investigates systematically reducing network complexity through top-down design. Our work builds on top of, but goes fundamentally beyond, prior research on network complexity [7], [10], [33], which focused on complexity measurement, quantification and modeling. Furthermore, this work is the first to reveal a fundamental limitation of the commonly accepted “divide-and-conquer” design approach [34], [35] in containing network complexity. We consider this insight a major advance in the state of the art in top-down network design.

For future work, we will seek to (i) extend the unified access control design framework by modeling also the task of routing

design, (ii) evaluate the framework on additional configuration datasets, (iii) validate the generality of the integrated design methodology by applying it to other operation objectives such as QoS and resiliency, as well as to the “one big switch” model of SDN, (iv) incorporate other types of network complexity metrics, including those with a higher level semantics about network behaviors than the configuration driven metrics, and last but not the least, (v) investigate how the presented approach, currently targeting new networks, can be adapted to support evolving and redesigning existing networks.

REFERENCES

- [1] Cisco IP Solution Center. [Online]. Available: <http://www.cisco.com/en/US/products/sw/netmgtsw/ps4748/index.html>, accessed Feb. 1, 2015.
- [2] Intelliden. [Online]. Available: <http://www.intelliden.com/>.
- [3] Opsware. [Online]. Available: <http://www.opsware.com/>.
- [4] Voyence. [Online]. Available: <http://www.voyence.com/>.
- [5] J. van Rantwijk. (2008). *Maximum Weighted Matching*. [Online]. Available: <http://jorisvr.nl/maximummatching.html>.
- [6] M. Behringer and G. Huston. (2013). *A Framework for Defining Network Complexity, Internet Draft (Work in Progress)*. [Online]. Available: <http://tools.ietf.org/html/draft-irtf-ncrg-complexity-framework-00>.
- [7] T. Benson, A. Akella, and D. Maltz, “Unraveling the complexity of network management,” in *Proc. USENIX NSDI*, 2009, pp. 335–348.
- [8] T. Benson, A. Akella, and D. A. Maltz, “Mining policies from enterprise network configuration,” in *Proc. 9th ACM SIGCOMM Conf. Internet Meas. Conf.*, 2009, pp. 136–142.
- [9] M. Caesar *et al.*, “Design and implementation of a routing control platform,” in *Proc. USENIX NSDI*, 2005, pp. 15–28.
- [10] B.-G. Chun, S. Ratnasamy, and E. Kohler, “NetComplex: A complexity metric for networked system designs,” in *Proc. USENIX NSDI*, 2008, pp. 393–406.
- [11] *Catalyst 2950 Desktop Switch Software Configuration Guide*, Cisco, San Jose, CA, USA, 2002.
- [12] Distributed Management Task Force, Inc. [Online]. Available: <http://www.dmtf.org>.
- [13] W. Enck *et al.*, “Configuration management at massive scale: System design and experience,” in *Proc. USENIX*, 2007, Art. no. 6.
- [14] N. Feamster and H. Balakrishnan, “Detecting BGP configuration faults with static analysis,” in *Proc. USENIX NSDI*, 2005, pp. 43–56.
- [15] N. Foster *et al.*, “Frenetic: A network programming language,” in *Proc. ACM SIGPLAN Int. Conf. Funct. Program.*, 2011, pp. 279–291.
- [16] J. Gottlieb, A. Greenberg, J. Rexford, and J. Wang, “Automated provisioning of BGP customers,” *IEEE Netw.*, vol. 17, no. 6, pp. 44–55, Nov./Dec. 2003.
- [17] T. G. Griffin and J. L. Sobrinho, “Metarouting,” in *Proc. ACM SIGCOMM*, 2005, pp. 1–12.
- [18] K. He, Y. Wang, X. Wang, W. Meng, and B. Liu, “GreenVLAN: An energy-efficient approach for VLAN design,” in *Proc. Int. Conf. Comput., Netw. Commun. (ICNC)*, Jan./Feb. 2012, pp. 522–526.
- [19] Juniper Networks. (2008). *What’s Behind Network Downtime?* [Online]. Available: <http://www-935.ibm.com/services/tw/gts/pdf/200249.pdf>.
- [20] N. Kang, Z. Liu, J. Rexford, and D. Walker, “Optimizing the ‘one big switch’ abstraction in software-defined networks,” in *Proc. ACM CoNEXT*, 2013, pp. 13–24.
- [21] Z. Kerravala, “As the value of enterprise networks escalates, so does the need for configuration management,” Yankee Group, Boston, MA, USA, Tech. Rep., 2004.
- [22] A. R. Khakpour and A. X. Liu, “Quantifying and querying network reachability,” in *Proc. IEEE ICDCS*, Jun. 2010, pp. 817–826.
- [23] F. Le, G. G. Xie, D. Pei, J. Wang, and H. Zhang, “Shedding light on the glue logic of the Internet routing architecture,” in *Proc. ACM SIGCOMM*, 2008, pp. 39–50.
- [24] F. Li *et al.*, “CSS-VM: A centralized and semi-automatic system for VLAN management,” in *Proc. IFIP/IEEE Int. Symp. Integr. Netw. Manage. (IM)*, May 2013, pp. 623–629.
- [25] F. Li, J. Yang, C. An, J. Wu, and X. Wang, “Towards centralized and semi-automatic VLAN management,” *Int. J. Netw. Manage.*, vol. 25, no. 1, pp. 52–73, Jan./Feb. 2015.
- [26] A. X. Liu, E. Torng, and C. R. Meiners, “Firewall compressor: An algorithm for minimizing firewall policies,” in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 691–699.
- [27] C. Monsanto, J. Reich, N. Foster, J. Rexford, and D. Walker, “Composing software-defined networks,” in *Proc. USENIX NSDI*, 2013, pp. 1–14.
- [28] S. Narain, “Network configuration management via model finding,” in *Proc. LISA Conf.*, 2005, p. 15.
- [29] *Networking in the Era of Virtualization*, Nicira, Palo Alto, CA, USA, 2012.
- [30] P. Oppenheimer, *Top-Down Network Design*, 3rd ed. Indianapolis, IN, USA: Cisco Press, 2010.
- [31] J. Pescatore, “Taxonomy of software vulnerabilities,” Gartner Group, Stamford, CT, USA, Tech. Rep. G00117235, 2003.
- [32] R. Rastogi, Y. Breitbart, M. Garofalakis, and A. Kumar, “Optimal configuration of OSPF aggregates,” *IEEE/ACM Trans. Netw.*, vol. 11, no. 2, pp. 181–194, Apr. 2003.
- [33] X. Sun, S. G. Rao, and G. G. Xie, “Modeling complexity of enterprise routing design,” in *Proc. ACM CoNEXT*, 2012, pp. 85–96.
- [34] X. Sun, Y.-W. E. Sung, S. D. Krothapalli, and S. G. Rao, “A systematic approach for evolving VLAN designs,” in *Proc. IEEE INFOCOM*, Mar. 2010, pp. 1–9.
- [35] Y.-W. E. Sung, X. Sun, S. G. Rao, G. G. Xie, and D. A. Maltz, “Towards systematic design of enterprise networks,” *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 695–708, Jun. 2011.
- [36] V. Tabatabaee, B. Bhattacharjee, R. J. La, and M. A. Shayman, “Differentiated traffic engineering for QoS provisioning,” in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 2349–2359.
- [37] Wikipedia. *K-Means Clustering*. [Online]. Available: https://en.wikipedia.org/wiki/K-means_clustering, accessed Feb. 1, 2015.
- [38] G. G. Xie *et al.*, “On static reachability analysis of IP networks,” in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 2170–2183.



Xin Sun received the B.E. degree in computer engineering from the University of Science and Technology of China in 2005, and the Ph.D. degree in computer engineering from Purdue University in 2012. Since 2012, he has been an Assistant Professor with the School of Computing and Information Sciences, Florida International University, Miami, FL, USA. In 2014, he was a Visiting Researcher with the IBM T. J. Watson Research Center. His research interests are in computer networks and networked systems, with a current focus on network design and management, network complexity, and software-defined networking. He was a recipient of the NSF CRII Award in 2015.



Geoffrey G. Xie received the B.S. degree in computer science from Fudan University, China, and the Ph.D. degree in computer science from The University of Texas at Austin. He was a Visiting Scholar with the School of Computer Science, Carnegie Mellon University, from 2003 to 2004, and the Computer Laboratory, University of Cambridge, U.K., from 2010 to 2011. He is currently a Professor and the Associate Chair with the Computer Science Department, Naval Postgraduate School. He has authored over 70 articles in various areas of networking. His current research interests include formal network analysis, routing design and theories, cloud security, and abstraction driven design of enterprise networks. He won the best paper award at the 2007 IEEE ICNP Conference and the 2011 IEEE Fred W. Ellersick Award. He was the Co-Chair of the ACM SIGCOMM Internet Network Management Workshop in 2007.