

# Characterization Study of VLANs in a Campus Network

Prashant Garimella, Yu-Wei Sung, Nan Zhang, Sanjay Rao  
Purdue University

## Abstract

In this paper we present a study characterizing VLAN usage in a large-sized campus network. Despite their extensive prevalence in enterprise and campus networks, the usage of VLANs has received little systematic treatment in the research community. Our study is conducted using a white-box approach, involving data such as router configuration files obtained from network operators, and through iterative interactions with them. Our study shows that the use of virtualization is prevalent to enable users belonging to physically disparate locations to be treated as a group. We demonstrate and characterize the performance inefficiencies resulting from virtualization. We show the inefficiencies are exacerbated by sub-optimal placement policies. We also discuss potential sources of errors that may arise with configuration of VLANs, and demonstrate their prevalence in real configurations. We believe these results are a key step towards gaining deeper insights into operational practices in enterprise and campus networks, and the design of abstractions to simplify management.

## 1 Introduction

In recent years, researchers have advocated the need for abstractions that model the fundamental design intent of a network manager’s actions, and capture the ultimate network-wide performance, security, manageability and resilience objectives of the designer [4]. While there has been tremendous attention and progress towards the design of network-wide abstractions in certain domains, most notably BGP [1, 5, 2], surprisingly little attention has been paid to the management of enterprise and campus networks. Despite their critical importance, and their striking differences and diversity compared to carrier networks, there is little systematic understanding about these networks in the community.

In this paper, we take a step towards addressing this by conducting a “bottom-up” study of network designs used in a real operational campus network. We believe the deeper understanding so obtained is a necessary first step to capturing the goals operators have for their networks, and can guide abstraction design. Our studies are based on unique “white-box” methodologies that involve access to data such as router configuration files obtained with the support of network managers, and iterative interactions with operators. Such white-box

methodologies are in contrast to black-box approaches used to “infer” characteristics about the network with limited operator support, and which are widely used in the community (e.g. [8, 11, 10]). While such a white-box methodology has been employed before (e.g. [9]), its usage is rare, and takes significant effort to boot-strap.

We focus in this paper on characterizing the usage of Virtual LANs (VLANs) in an operational campus network. VLANs are extensively used in campus and enterprise networks, pose significant challenges to network managers, but receive almost no attention in the research community. VLANs are often used to address groups of users as a single unit to ease management, even though they may be spread over physically disparate locations and not connected to the same routers or switches. For instance, a policy in an enterprise may permit access to all sales personnel alone, and it may be desirable to ensure these users receive IP addresses from the same subnet so that IP routing policies and packet filters can be applied to them as a group. Configuring VLANs is a manual process and represents an activity that managers spend much of their time on.

Based on discussions with operators, we abstract and expose key issues that must be considered in designing a VLAN architecture. We conduct a study of the use of VLANs in Purdue University. The network consists of about 200 routers, 1300 switches, and a few hundred VLANs. While campus networks are different than enterprises in general, the size of the network, availability of data, and the extensive use of VLANs makes the Purdue network a good data-point, and many of the issues can be generalized to other networks. We show that the usage of VLANs is prevalent, and virtualization is often used to span disparate physical locations. We demonstrate and characterize the performance inefficiencies resulting from such virtualization. We show the inefficiencies are exacerbated by sub-optimal placement policies. We also discuss potential sources of errors that may arise with configuration of VLANs, and demonstrate their prevalence in real configurations.

Although the complexity of VLANs has been recognized by other researchers [6], to our knowledge, this is the first systematic exposition of issues in designing VLANs, and the first empirical characterization of VLANs in a real network. We believe these results are a key step towards gaining deeper insights into oper-




Figure 1: Example VLAN setup. Communication between different VLANs is routed through designated routers.

S1	S
interface FastEthernet0/1 switchport mode access switchport access vlan 1 !	interface FastEthernet0/1 switchport mode trunk switchport trunk allowed vlan 1,2 !
interface FastEthernet0/2 switchport mode trunk switchport trunk allowed vlan 1,2 !	interface FastEthernet0/2 switchport mode trunk switchport trunk allowed vlan 1,3 !
R1	R2
interface Vlan1 description subnet 192.168.1.0/24 ip address 192.168.1.1 255.255.255.0 !	interface Vlan2 description subnet 192.168.2.0/24 ip address 192.168.2.1 255.255.255.0 !

Figure 2: VLAN configurations for devices in Figure 1.

ational practices in configuring VLANs and managing enterprises in general. These insights in turn can enable the design of abstractions to simplify management, and inform the design of clean-slate architectures.

## 2 VLAN Design

In this section, we present the background on VLANs, and considerations in designing a VLAN architecture. These considerations have arisen out of our discussions with operators, and we believe highlighting them forms a key contribution of this paper.

### 2.1 Background

Consider Figure 1. S, S1-S3 are switches, and R1 and R2 are routers. Hosts H1 and H3 belong to VLAN 1, and hosts H2 and H4 belong to VLAN 2. The relevant configuration snippet of switch S and S1 is shown in Figure 2. The link between S1 and H1 is configured as an *access* link and only traffic of VLAN 1 is forwarded on that link. The link between S1 and S is configured as a *trunk* link. A trunk link may carry traffic corresponding to multiple VLANs, and the list of VLANs allowed on that link must be explicitly configured on both ends. In the example, the trunk is configured to allow traffic corresponding to VLANs 1 and 2, as there are hosts on both sides of the link belonging to each VLAN.

Each VLAN is designated as public, or private, and is usually associated with an IP subnet. Hosts in private VLANs can only communicate with other hosts within the VLAN. Therefore, private VLAN numbers

can be reused to represent multiple different VLANs. Each public VLAN is assigned with what we term a *designated router* for that VLAN. In Figure 1, R1 and R2 are respectively the designated routers for VLAN 1 and VLAN 2. Figure 2 shows the relevant configuration snippet in R1 and R2 to indicate that this is the case and their associated subnets. When a host in a VLAN communicates with a host outside, the designated router is the first (last) router for outgoing (incoming) packets. Note that while we have separated switches and routers completely, in practice, some devices act as both routers and switches, while others act as switches alone.

### 2.2 Design Issues

Our discussions with operators lead us to abstract key issues that must be considered when designing VLANs: **Performance inefficiencies:** While virtualization simplifies management, it introduces inefficiencies. To understand this, for hosts in two different VLANs, consider (i) the shortest physical path, and (ii) the path that data actually flows between them. For example, in Figure 1, the shortest physical path between H1 and H2 is simply  $H1 - S1 - H2$ , as both hosts are attached to the same switch. The path along which data flows is obtained by considering the IP-level path, and for each of its IP hops, considering the shortest physical path. The IP level path between H1 and H2 is:  $H1 - R1 \dots R2 - H2$  ( $R1 \dots R2$  denoting there could be other routers in the path), and the path of data flow is as shown in the figure. Note that R1 acts as a router in outgoing direction, but as a switch in the return direction.

Using the substantially longer paths for data flows may involve longer delays, redundant transmission, and loops. While the concern may not appear critical as links are usually under-utilized, operators prefer to avoid inefficiencies in order to provision against new applications with unexpected traffic patterns, or worm out-breaks. These concerns are particularly important in key links, for instance those that connect a building to the core. Further, the longer paths increase the likelihood of failures, and complicate performance and failure diagnosis. For example, in Figure 1, communication between H1 and H2 may be affected by the failure of any of the devices along the data flow path, even though a physical path exists. If H1 and H2 were in a building in campus, and other devices located in external buildings, communication could be disrupted by issues such as power failure in external buildings.

**Placement of designated router:** While inefficiencies are inherent to VLANs, the extent of inefficiency is impacted by the placement of the designated router for the VLAN. The smaller the “distance” between a host and the *designated router* of its VLAN, the lower the inefficiency is. For example, in Figure 1, the inefficiencies of communication between H1 and H2 would

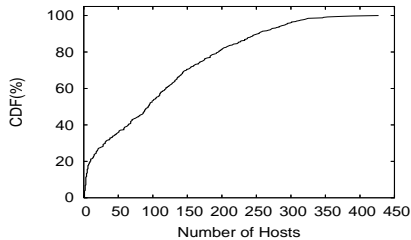


Figure 3: Number of hosts in a VLAN.




Figure 4: Number of buildings spanned by a VLAN.




Figure 5: Number of building-sets & core-sets spanned by a VLAN.

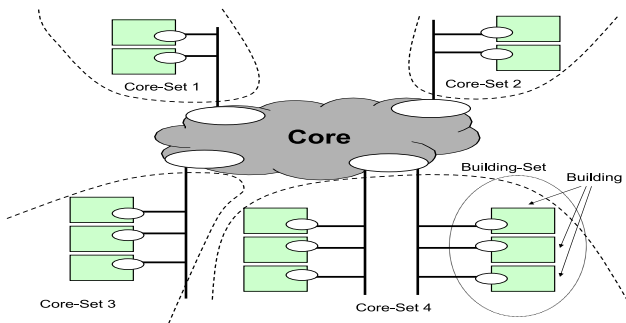


Figure 6: Outline of Purdue's campus network topology.

be minimized if R1 were chosen as the first hop router of VLAN 2 instead of R2. However, this could lead to higher inefficiencies for communication involving host H4, which also belongs to VLAN 2. Thus, an ideal placement strategy must consider the overall span of a VLAN, and must choose a designated router as close as possible to the majority of hosts in that VLAN. Other factors that could influence the decision include the traffic patterns of hosts - for instance if hosts in a VLAN tend to communicate with certain servers more often, a placement strategy that places the designated router closer to those servers may be preferable.

**Configuring allowed lists:** A key aspect in VLAN configuration is that both ends of every *trunk* link need to be explicitly configured with a list of VLANs allowed on that link. In general, if two hosts of a VLAN are located on both sides of a trunk link, traffic of that VLAN must be permitted on the link to ensure they can communicate. In Figure 1, H1 and H3 belong to VLAN 1, and hence both end-points of trunk links S-S1, and S-S2 must be configured to permit traffic from VLAN 1. Figure 2 shows the appropriate configuration snippet for switch S and S1. While it is important to permit VLANs on a link when required to ensure connectivity, it is also important to avoid permitting VLANs when not required. Since a switch forwards broadcast traffic for a VLAN to all ports allowing that VLAN, constraining the VLANs permitted on a trunk link minimizes unnecessary propagation of broadcast traffic. For example, in Figure 1, since all hosts of VLAN 2 are clustered on one side of trunk link S-S2, the link should not permit traffic belonging to VLAN 2, to prevent its broadcast traffic from leaking through to devices attached to S2. This

may become particularly important to protect against hosts that accidentally (or even deliberately) introduce large volumes of broadcast traffic into the network. As the number and scale of VLANs increase, proper configuration of the allowed lists on each trunk link becomes extremely difficult.

### 3 Operational Network Study

We study the network with a view to understanding the following questions: (i) how prevalent is the usage of VLANs to treat hosts in physically disparate locations as one unit?; (ii) how significant are the performance inefficiencies induced by VLANs?; (iii) how reasonable are the heuristics used for placement of designated routers?; and (iv) are there errors in the configuration of allowed lists in VLAN trunks? Our study was conducted on the Purdue network which consists of about 200 routers, 1300 switches, and a few hundred VLANs.

#### 3.1 Prevalence of virtualization

Figure 6 depicts a conceptual outline of the Purdue campus topology. A small bunch of routers form the core of the topology. Typically, each building has a router with a link to the core. This router connects all hosts in that building to the rest of the campus network. We call this router the *primary router* of a building. We define buildings (and their hosts) that connect to the same core router as belonging to the same *core-set*. A core-set may comprise one or more *building-sets* - these refer to buildings that not only connect to the same core-router but also tend to have related characteristics (e.g. corresponding to different engineering departments).

Figure 3 plots a CDF of the number of hosts in each VLAN. A point  $(x, y)$  indicates that  $y\%$  of VLANs have  $x$  hosts or fewer. Nearly half of the VLANs involve over a hundred hosts. Figure 4 considers how many buildings are spanned by the hosts in a VLAN. A point  $(x, y)$  indicates that a fraction  $y$  of VLANs span  $x$  buildings or fewer. While 50% of the VLANs span only one building, about 10% of the VLANs span more than 5 buildings, and the largest VLAN spans over 60 buildings!

We consider the span of VLANs further to capture whether they tend to be clustered in the same region of




Figure 7: Routing inefficiency between two hosts located in the same building but from different VLANs.

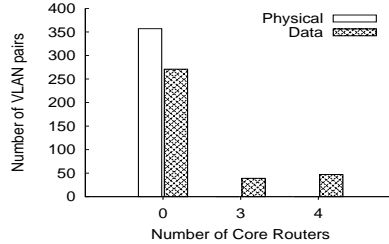


Figure 8: Number of core routers traversed for inter-VLAN communications.

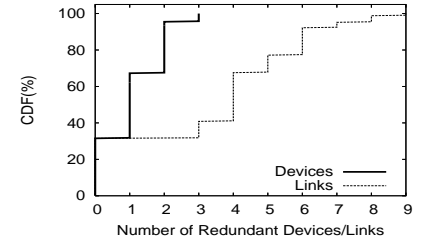


Figure 9: Number of redundant devices & links traversed for inter-VLAN communications.

campus. Figure 5 has two curves. The lower curve plots a CDF of the number of building-sets spanned by the VLAN, while the top curve plots a CDF of the number of core-sets spanned by the VLANs. Many VLANs do tend to be clustered around similar parts of campus. For instance, about 70% of VLANs span a single building-set, and 80% of VLANs span just one core-set. However, a small fraction of the VLANs do span several building-sets and core-sets, indicating they are distributed over diverse locations on campus.

We discussed further with operators to understand typical causes for the span of VLANs. One of the VLANs that span several buildings contains hosts in all classrooms on campus. Another VLAN includes hosts in all conference rooms across campus. Typically all of these VLANs are managed by the same administrative sub-unit distinct than the main campus operators, and use of VLANs simplifies the task of allocating IP address blocks to the sub-units.

### 3.2 Performance inefficiencies

We obtained raw topology information from network managers that included information regarding all physical devices, and the links connecting them, which enabled us to find the shortest path between any two devices. In addition to physical topology, we also obtained the routing tables of core routers, and an understanding of the structure of routing. This enabled us to determine the path between any pair of IP addresses, and in turn the path that data flows between the addresses.

To characterize the performance inefficiencies, we use the *routing inefficiency* metric, defined as the ratio of the number of hops on the path that data flows between two hosts to the number of hops on the shortest physical path between them. A building is considered, and for each pair of distinct VLANs with hosts in that building, representative hosts are selected, and the routing inefficiency is computed for that pair of hosts. Figure 7 shows the CDF of routing inefficiencies for 3 buildings. Each curve corresponds to a building, the X-Axis is the routing inefficiency, and the Y-Axis shows the fraction of VLAN pairs for which the inefficiency is less than a particular value. For all buildings, the inefficiency is

significant, and greater than 4 for 12% of the pairs.

While the inefficiencies are high, we look at metrics that can better capture the practical significance of the longer data paths with VLANs. Figure 8 shows the number of core routers traversed in the physical and data path for Building1. Intuitively, it is desirable to minimize traffic in the core, and the more core routers are traversed, the more significant the inefficiency concerns are. For each pair of VLANs in Building1, representative hosts in Building1 belonging to those VLANs are taken, and the number of core routers traversed in the physical path, and path of data flow is considered. Each bar corresponds to the fraction of pairs for which a certain number of core routers are traversed. When physical paths are considered, all pairs involve no core router. When paths involving data flow are considered; however, about 11% of pairs involve 3 core routers, and about 13% of pairs involve 4 core routers.

An offshoot of the inefficiency concerns involves the formation of loops, where the same data traverses the same device/link multiple times. Figure 9 studies this further. For each pair of distinct VLANs in Building1, the path of data-flow between two representative hosts is considered, and we examine (i) the number of devices (switches/routers), and (ii) the number of links that appear multiple times on the path. Note that a link refers not just to connections between switches and routers, but also includes connections to optical interconnect boxes, and patch-panels. The CDF of the number of redundant devices, and links across VLAN pairs is plotted. About 70% of pairs involve at least one device that appears redundantly in the path, and about 70% of the pairs involve at least one link that appears redundantly on the path. Some paths may involve two switches/routers that appear multiple times, and as many as 7 or 8 links that appear twice.

While our results so far have considered VLAN pairs with hosts located in the same building, we also considered VLAN pairs with hosts located in two different buildings. While the overall results are similar we found interesting cases where the same link can be traversed three times. In particular, there were two buildings B1 and B2, both of which had hosts belonging to VLANs

VLAN Span	Total VLANs	Satisfy BLDG MOSTHOSTS rule	# of Cases	Satisfy CORESET MOSTHOSTS rule	# of Cases
Single Building	149	Yes	133	Yes	147
		No	16	No	2
				Unknown	1
Exactly 2 Buildings	60	Yes	40	Yes	57
		No	20	No	1
				Unknown	2
> 2 Buildings	96	Yes	40	Yes	70
		No	56	No	19
				Unknown	7

Table 1: Characteristic of designated router placement.

V1 and V2, with the designated routers of the VLANs R1 and R2 lying in buildings B1 and B2 respectively. The inefficiencies resulted when a host X in VLAN V1 and Building B2, communicated with a host Y in VLAN V2 and Building B1.

### 3.3 Placement of designated routers

While the previous section has shown performance inefficiencies with VLANs, in this section we investigate this further to see if the inefficiencies could have been minimized by more careful placement of designated routers. Our analysis was conducted using configuration files of all switches and routers in the campus. The information helped identify the designated router of each VLAN. Our analysis excluded private VLANs. In cases where a VLAN had multiple designated routers, with a primary and multiple stand-by routers, we considered the location of the primary router.

We employed two heuristics to evaluate the placement of the designated routers. These heuristics were rules of thumb used by operators in making placement decisions:

- *BLDG-MOSTHOSTS*: We considered whether the designated router was placed in the building that had the most hosts in that VLAN. If the VLAN was entirely contained in 1 building, this simply translates to whether the designated router is in that building.
- *CORESET-MOSTHOSTS*: While the rule above was typically employed by operators if a building had 50% or more of the hosts in a VLAN, in cases where a VLAN was evenly spread across multiple buildings, the designated router was placed in the same core-set (Section 3.1) as the core-set with the most hosts in that VLAN.

Table 1 summarizes our analysis to examine whether the placement of designated routers conformed to the two rules above. There were 149 VLANs contained entirely in a single building, 60 VLANs that spanned two buildings, and 96 VLANs which covered more than two buildings. Overall 89% (133/149) of the single building VLANs, 67% (40/60) of the two-building VLANs, and 42% (40/96) of multi-building VLANs conformed to the *BLDG-MOSTHOSTS* rule. Further analysis indicated there were striking cases of sub-optimal placements. Among the 20 2-building VLANs that did not conform, there were 3 cases where the designated router

was placed in a building with less than 10% of the hosts, and 8 cases where the designated router was in neither of the buildings. Among the 56 multi-building VLANs that did not conform, there were 34 cases where the building with most hosts had 70% or more of all hosts in the VLAN, and 6 cases where the building with most hosts had 90% or more of all hosts in the VLAN.

Table 1 also presents results regarding whether the placements conform to the *CORESET-MOSTHOSTS* rule. Since this rule represents a more relaxed interpretation of good placement, the heuristics do show better results. Practically all single and 2-building VLANs conform to this rule. However, when VLANs that span multiple buildings are considered, only 73% (70/96) of the VLANs conform to the rule. There are 19 VLANs have sub-optimal placement. The non-conformance was particularly striking in 5 of these cases, where the designated router was located in a core-set with no known hosts in that VLAN, despite the core-set with most hosts has over 70% of the hosts in that VLAN.

We had further discussions with operators to understand why sub-optimal placements occurred. In some cases the explanation was straight-forward - a few buildings did not have routers, and the designated router had to be selected from other available buildings. In a few cases, the placement was deliberate based on the known traffic patterns of hosts - for example, if hosts in that VLAN frequently accessed a small set of servers, the designated router could be placed close to the servers. However, there were other cases where the sub-optimal placement was an artifact of changes to the network - as hosts were added and removed to VLANs, previous choices of reasonable placement were no longer appropriate. Finally, there were cases which just resulted from operator oversight, pointing to the manual, error-prone, and complex nature of the placement problem.

### 3.4 Configuration of allowed VLANs

While our results have focused on performance inefficiencies with VLANs, and the role of sub-optimal placement, we next consider errors that arise due to misconfiguration of the VLANs permitted on trunk links.

We considered the following types of errors:

- *Missing VLANs*: Here, a VLAN that should be specified in the allowed list of a trunk link is omitted. This may result in disconnection between hosts belonging to that VLAN located on different sides of the trunk link.
- *Unnecessary VLANs*: Here, a VLAN is unnecessarily specified on a trunk link. It is unnecessary in that all hosts in the VLAN are located on one side of the trunk link, and can communicate in the absence of that link. This misconfiguration may lead to superfluous broadcast traffic being flooded on that link.

To study the prevalence of these errors, we parse the configurations to extract the list of VLANs allowed on

Missing VLANs	Tot Bldg	Number of VLANs	
		Total	Hosts affected
Misconfigurations	5	8	Avg:2.25,Max:4
Nested building	4	8	0
Unnecessary VLANs	Tot Bldg	Number of Errors	
		Extra out	Extra in
Misconfigurations	119	53	6574
Desig. router outside	6	-	-

Table 2: Causes of allowed list misconfiguration.

all trunk links that connect the primary router of a building to the core. Table 2 summarizes our findings. Among 131 configuration files analyzed corresponding to buildings with a direct connection to a core router, only 5 had errors corresponding to missing VLANs. Further, 8 VLANs were affected, and each error impacted (disconnected) 2.25 hosts on average, and a maximum of 4. It is reasonable that these errors are small, since these would lead to complaints from hosts (users) that are disconnected from the network. In fact, our discussions with operators reveal there have been multiple real incidents of missing VLAN errors in configurations that have been fixed in response to user complaints. However, when unnecessary VLANs are considered, 119 of 131 configurations contain such errors. Further, there were 53 cases where VLANs entirely located in a building were unnecessarily specified on a trunk leading to superfluous broadcast traffic exiting the building, and 6574 cases where VLANs entirely located outside the building were specified in trunk links potentially leading to unnecessary broadcast traffic entering a building. Overall these results point to the complexity, and the manual and error-prone nature of the problem of configuring VLANs on trunk links.

## 4 Summary and Discussion

While VLANs are extensively used in practice, and represent a critical and time-consuming activity in enterprise network management, they are poorly understood and have received little attention from the research community thus far. In this paper we demystify some of the black art surrounding the management of VLANs, expose issues that operators grapple with, and provide a framework to understand the issues involved.

We have conducted the first and most extensive evaluation of an actual VLAN design in an operational campus network using a white-box methodology. While campus networks are distinct than enterprises in general, the size of the network we consider, the availability of data, and the extensive use of VLANs makes our study a great starting point. Further, we believe many of the issues are general to other networks.

Our evaluations show that the use of virtualization to span physically disparate locations is widely prevalent to simplify coordination between the campus administrators, and other administrative sub-units. Virtualization results in significant performance inefficiencies.

Minimizing inefficiencies, requires operators to carefully place designated routers. Despite significant effort spent in the process, network evolution and operator oversight leads to sub-optimal placement. Configuring VLANs permitted on trunks is error-prone and complex, leads to hosts being disconnected, and result in unnecessary broadcast traffic. While the former class of errors is often detected during operation and fixed, the latter class is quite prevalent, leaving a network susceptible to hosts that may accidentally or even deliberately introduce large volumes of broadcast traffic into the network.

We believe the understanding and insights in this paper opens the door for the research community to further engage in the area. First, VLANs are an excellent case study for the design of network-wide abstractions. Ideally, while a manager should specify high-level goals such as which hosts must be treated as a unit, and desired performance targets, low-level choices such as designated router placement and trunk configuration must be automatically generated. Second, VLANs are an interesting case-study for change management - addition of a host to the network involves network-wide dependencies in that in not only involves modification of the configuration of the switch it is attached to, but also modifications to the other switches and elements (e.g. allowed VLAN lists), and may require periodic reevaluation of location of first-hop router. Third, VLANs present new issues for failure and performance diagnosis. Further, they may serve as interesting application areas for fault diagnosis techniques such as [7]. Finally, our insights and data can complement and inform the design of clean-slate architectures [4, 3, 6], which have been partially motivated by the complexity of VLANs.

## References

- [1] C. Alaettinoglu, D. Meyer, and J. Schmitz. Application of routing policy specification language, 1997.
- [2] H. Boehm, A. Feldmann, O. Maennel, C. Reiser, and R. Volk. Network-wide inter-domain routing policies: Design and realization. Apr. 2005.
- [3] M. Casado, T. Garfinkel, A. Akella, M. Freedman, D. Boneh, N. McKeown, and S. Shenker. SANE: A protection architecture for enterprise networks. In *Usenix Security*, Aug. 2006.
- [4] A. Greenberg, G. Hjalmytsson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang. A clean slate 4D approach to network control and management. *ACM Computer Communication Review*, October 2005.
- [5] T. G. Griffin and J. L. Sobrinho. Metarouting. In *Proc. ACM SIGCOMM*, Aug. 2005.
- [6] C. Kim and J. Rexford. Revisiting ethernet: Plug-and-play made scalable and efficient. In *Proc. IEEE LANMAN Workshop*, 2007.
- [7] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren. Ip fault localization via risk modeling. In *Proc. Networked Systems Design and Implementation*, 2005.
- [8] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson. User-level Internet Path Diagnosis. In *Proc. of SOSP*, 2003.
- [9] D. Maltz, G. Xie, J. Zhan, H. Zhang, G. Hjalmytsson, and A. Greenberg. Routing design in operational networks: A look from the inside. In *Proc. ACM SIGCOMM*, Aug. 2004.
- [10] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel. In *Proc. of ACM SIGCOMM*, 2002.
- [11] M. Zhang, C. Zhang, V. Pai, L. Peterson, and R. Wang. Planeteer: Internet Path Failure Monitoring and Characterization in Wide-Area Services. In *Proc. of OSDI*, 2004.