

DOCSDN: Dynamic and Optimal Configuration of Software-Defined Networks

Timothy Curry, Devon Callahan, Benjamin Fuller and Laurent Michel

Department of Computer Science and Engineering

University of Connecticut Storrs, CT USA

{timothy.curry,devon.callahan,benjamin.fuller,laurent.michel}@uconn.edu

Abstract—Networks are designed with functionality, security, performance, and cost in mind. Tools exist to check or optimize individual properties of a network. These properties may conflict, so it is not always possible to run these tools in series to find a configuration that meets all requirements. This leads to network administrators manually searching for a configuration.

This need not be the case. In this paper, we introduce a layered framework for optimizing network configuration for functional and security requirements. Our framework is able to output configurations that meet reachability, bandwidth, and risk requirements. Each layer of our framework optimizes over a single property. A lower layer can constrain the search problem of a higher layer allowing the framework to converge on a joint solution.

Our approach has the most promise for software-defined networks which can easily reconfigure their logical configuration. Our approach is validated with experiments over the fat tree topology, which is commonly used in data center networks. Search terminates in between 1-5 minutes in experiments. Thus, our solution can propose new configurations for short term events such as defending against a focused network attack.

Index Terms—network security, reachability, risk assessment, optimization, software defined networks.

I. INTRODUCTION

Network configuration is a crucial task in any enterprise. Administrators balance functionality, performance, security, cost and other industry specific requirements. The resulting configuration is subject to periodic analysis and redesign due to red team recommendations, emerging threats, and changing priorities. Tools assist administrators with this complex task: existing work assesses network reachability [1], wireless conflicts [2], network security risk [3], [4], and load balancing [5], [6]. These tools assess the quality of a potential configuration. Unfortunately, current tools suffer from three limitations:

- 1) Most tools assess whether a single property is satisfied, making no recommendation if the property is not satisfied. This leaves IT personnel with the task of deciding how to change the network.
- 2) Networks are assessed with respect to an individual goal at a time. This means a change to satisfy a single property may break another property. There is no guidance for personnel on how to design a network that meets the complex and often conflicting network requirements.
- 3) These tools do not react to changing external information such as the publication of a new security vulnerability.

a) Our Contribution: This work introduces a new optimization framework that finds network configurations that satisfy multiple (conflicting) requirements. We focus on data center networks (DCN) that use software defined networking (SDN). Background on these settings is in Section II. Our framework is called DOCSDN (Dynamic and Optimal Configuration of Software-Defined Networks).

DOCSDN searches for network configurations that simultaneously satisfy multiple properties. DOCSDN is organized into layers that consider different properties. The core of DOCSDN is a multistage optimization that decouples search on “orthogonal” concerns. The majority of the technical work is to effectively separate concerns so the optimization problems remain tractable. Our framework is designed to continually produce network configurations based on changing requirements and threats. It frees IT personnel from the complex question of how to satisfy multiple requirements and can quickly incorporate new threat information.

DOCSDN focuses on achieving functional requirements (such as network reachability and flow satisfaction) and limiting security risk (such as isolating high risk nodes and nodes under denial of service attack). Naturally, other layers such as performance or cost can be incorporated. The search for a good configuration could be organized in many ways. State-of-the-art approaches assess different properties in isolation, frustrating search for a solution that satisfies all requirements. Ideally, a framework should search for a configuration that simultaneously satisfies all requirements. This extreme is unlikely to be tractable on all but the smallest networks. DOCSDN mediates between these approaches separating the functional and security search problems but introducing a feedback loop between the two search problems based on *cuts*.

In the proposed organization the functional layer is “above” the security layer. Through the feedback loop, the security layer describes a problematic part of the network to the functional layer. The functional layer then refines its model and searches for a functional configuration that satisfies an additional *constraint*. This has the effect of blocking the problematic part of the configuration. Currently, the feedback signal is a pair of nodes that should not be proximate in the network. After multiple iterations the two layers jointly produce a solution that optimizes the SDN configuration both with respect to functionality and security risks.

DOCSDN provides solutions of improving quality before

the final solution. Thus, the network can be reconfigured once the objective improves on the current configuration by a large enough amount (to justify the cost/impact of reconfiguration).

While the underlying optimization problems are NP-hard, optimization technology has seen tremendous advances in performance during the past few decades. Since 1991, mathematical programming solvers have delivered speedups of 11 orders of magnitude [7], [8]. The advent of hybrid techniques such as Benders decomposition [9]–[12] and column generation [13]–[15] (aka, Dantzig-Wolfe decomposition [16]) made it possible to harness huge problems thanks to on-demand generation of macroscopic variables and the dynamic addition of critical constraints. Large Neighborhood Search [17] further contributed to delivering high-quality solution within constrained time budgets.

These techniques are beginning to see adoption in network security. Yu et al. recently applied stochastic optimization with Bender’s decomposition to assess network risk under uncertainty for IoT devices [4]. They used Bender’s decomposition on a scenario-based stochastic optimization model to produce a master problem that chooses a deployment plan while slaves are concerned with *choosing* the optimal nodes to serve the demands in individual scenarios. In comparison, our approach addresses *both functional and security requirements*. It relies on Bender’s cut from the security layer (slave) to rule out vulnerable functional plans whose routing paths fail to adequately minimize risks and maximize served clients. We now briefly describe the framework (a formal description is in Section III) and present an illustrative example.

b) Overview of DOCSDN: Figure 1 presents an overview of the framework. The functional layer takes as input a *Functional Model* that describes the network including the physical topology, link speeds, the allowable communication patterns and the demand requirements. Network reachability begins with a priming procedure that generates the k -least cost paths to the optimizer for each source/destination pair in the demand requirements. The objective for the functional layer is to find a logical topology (a collection of routed paths) that maximizes the number of good flows that are served. The program is formulated as quadratic binary program (*QBP*). The solution as determined by the functional layer is passed to the security layer.

The output of the functional layer and a *security model* are the input for the security layer. The current configuration is fed to a module that uses risk assessments for the individual network devices (obtained for example using a vulnerability database) to assess the overall risk of the entire configuration. In our current implementation this risk calculation is based on a simple risk propagation model where a path’s risk is based on the risk of nodes on the path and close to the source and sink. The security layer can deploy firewalls and deep packet inspection as network defenses. Since these mechanisms affect route capacity, the security layer has a dual objective function: 1) maximizing the functional objective and 2) minimizing security risk. The security objective is formulated as a mixed integer program (*MIP*). When the security search completes, it

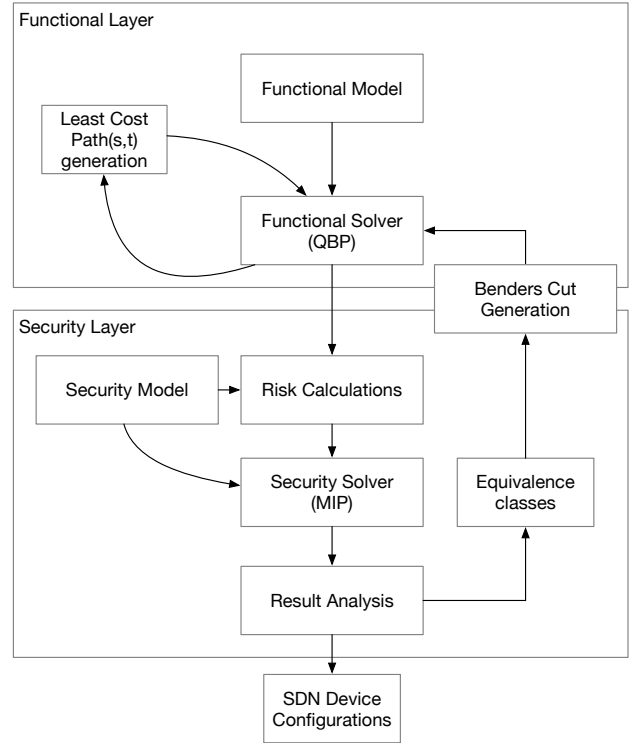


Fig. 1. DOCSDN Framework. A layered decomposition that breaks down configuration synthesis into functional and security layers.

proposes nodes to the functional layer that should be separated. As an example, a high value node with low risk may be placed in a different (virtual) LAN than a high risk node. These *Benders cuts* are designed to entice a better logical topology from a subsequent iteration in the functional layer. This feedback loop between the two layers can iterate multiple times. When no further cuts are available, the overall output is a set of configuration rules.

c) An example configuration: This section describes an application of our framework to automatically respond to a distributed denial of service (DDoS) attack. Current DDoS attacks demonstrate peak volume of 1 Tbps [18]. Many DDoS defense techniques require changes to the network behavior by rate limiting, filtering, or reconfiguring the network (see [19]–[23]). Recent techniques [24] leverage SDNs to react to DDoS attacks in a dynamic and flexible manner. We show how such a response would work in our framework using a toy network illustrated in Figure 2. A more realistic network and the framework’s response are described in Section V. We stress that DDoS attacks are often short in timescale making human diagnosis and reaction costly or impractical. Consider a focused DDoS attack against a number of services in an enterprise but not the entirety of its publicly accessible address space. (The Great Cannon’s attack against GreatFire targeted two specific Github repositories [25].) We assume a service hosted by H_1 is targeted, while services on H_2, H_3 and H_4 are not.

Recall, the functional layer establishes a logical topology




Fig. 2. Toy network example with a single gateway device G_1 , two intermediate switches S_1 and S_2 and four hosts. We assume the switches are physically connected to all hosts.

(forwarding rules) while the security layer adds network defenses (packet inspection modules and firewall rules). We elide how the attack is detected and assume it increases the risk score for H_1 in the security model.

The first iteration The functional layer proposes a candidate configuration where G_1 routes all traffic intended for H_1 and H_2 to S_1 which then forwards the traffic and G_1 routes traffic intended for H_3 and H_4 to S_2 which then forwards the traffic. This is the first candidate solution presented to the security layer.

Since H_1 is high risk the security layer proposes a firewall at S_1 to block all port 80 traffic. This reduces risk at the cost of blocking all traffic to H_2 . Of course, in real firewalls more fine-grained rules are possible, this simplified example is meant to illustrate a case where collateral damage to the functional objective is necessary to achieve the security objective. Since traffic is being blocked to a node with low risk, the security layer asks the functional layer to separate H_1 and H_2 so H_2 does not suffer.

Repeated iterations The functional layer now has a constraint that H_1 and H_2 should not be collocated in the network. As such, it proposes a new configuration with H_1 and H_3 under S_1 and H_2 and H_4 under S_2 . This is then sent to the security layer. The security layer makes a similar assessment and proposes a firewall rule at S_2 , finds this recommendation hurts functionality and requests separation of H_1 and H_3 .

This process repeats with the functional layer proposing to collocate H_1 and H_4 . The security layer similarly asks to separate H_1 and H_4 . Finally, H_1 is segregated from all other nodes. This produces a configuration where H_1 is the only child of S_1 . Note that having H_2, H_3 and H_4 under a single switch may hurt performance but the effect is less than blocking traffic to one of the nodes entirely. DOCSDN can then output the candidate solution as high level SDN fragments (using a high-level language like Pyretic).

Recovery Importantly, when the DDoS abates, DOCSDN automatically reruns with a changed risk for H_1 , outputting a binary tree.

Organization The rest of the work is organized as follows: Section II provides background on our application and discusses related work, Section III describes our framework and accompanying optimization models, Section IV describes our experimental setup, Section V evaluates the framework and finally Section VI concludes.

II. BACKGROUND AND RELATED WORK

Data Center Networks (DCN) host, process and analyze data in financial, entertainment, medical, government and military sectors. The services provided by DCNs must be reliable, accurate and timely. Services provided by DCNs (and the corresponding traffic) are heterogeneous. The network must adapt to changing priorities and requirements while protecting from emerging threats. They scale to thousands of servers linked through various interconnects. Protocols used for these services are split roughly 60 percent web (HTTP/HTTPS) and 40 percent file storage (SMB/AFS) [26]. The interdependence of device configurations make modifying any single configuration difficult and possibly dangerous for network health. A seemingly simple update can cause significant collateral damage and unintended consequences.

Simultaneously, the network fabric is changing with the advent of Software Defined Networking (SDN) [27]. SDNs are flexible and programmable networks that can adapt to emergent functional or performance requirements. Openflow [28] is a common open source software stack. Researchers have proposed high-level languages and compilers [29]–[32] that bridge the semantic gap between network administrators and the configuration languages used by SDN devices. These languages focus on *compositional* and *parametric* SDN software modules that execute specific micro-functions (e.g., packet forwarding, dropping, routing, etc.). The use of a high level language is prompted by a desire to be able to *select*, *instantiate* and *compose* SDN modules with guarantees.

Our framework is intended to be modular and allow integration of prior work on evaluating network configurations. As such there is a breadth of relevant work. Due to space constraints we focus on the most relevant works. In the conclusion we elaborate on the characteristics needed to integrate a prior assessment tool into our framework (see Section VI).

Measuring Network Risk Known threats against computer systems are maintained by governments and industry. Common Vulnerabilities and Exposures (CVE) is a publicly available dictionary including an identifier and description of known vulnerabilities [33], CVE does not provide a severity score or priority ranking for vulnerabilities. The US National Vulnerability Database (NVD) [34] is provided by the US National Institute of Standards and Technology (NIST). The NVD augments the CVE, adding severity scores and impact ratings for vulnerabilities in the CVE.

There are many mechanisms for measuring the security risk on a network [35]–[39]. Lippmann et al. present a network security model which computes risk based on a list of the most current threats [40]. This model implements a cycle of observe network state, compute risk, prioritize risk, and mitigate the risk.

This loop is often codified into an *attack graph* [3], [41], [42]. Attack graphs try to model the most likely paths that an attacker could use to penetrate a network. Attack Graphs often leverage one or more of the aforementioned vulnerability assessment tools as input, combined with a network topology and device software configurations to generate the graph.

Current attack graph technologies provide recommendations to network administrators that effectively remove edges from the graph and trigger a re-evaluation of the utility for the attacker. To the best of our knowledge, current practice does not leverage network risk measurement into constraints used for the generation of new configurations.

Network Reachability The expansion of SDN has aided the applicability of formal verification to computer networks. Prior to SDN, the lack of clear separation between the data and control plain created an intractable problem when considering a network of any scale. Bounded model checking using SAT and SMT solvers [43], [44] can currently verify reachability properties in networks with several thousands of nodes.

Configuration Search Constraint Programming (CP) was introduced in the late 1980s [45] and is used for scheduling [46], routing, and configuration problems. Large-scale optimization problems are often decomposed including Benders [12] and Dantzig-Wolfe [16]. *Soft constraints* or Lagrangian relaxation are used for over-constrained problems or when the problem is too computationally expensive. Stochastic optimization techniques have been used for many applications in resilience [47], [48] and the underlying methodologies are a key part of this research. Prior work in configuration management with constraint programming [49], [50] focused on connectivity or security. We are not aware of any work that balances these two objectives in a meaningful way.

III. IMPLEMENTATION

Figure 1 outlines the overall structure of the DOCSDN framework. layer inter-connections as well as their internals. The functional layer uses a mathematical optimization model that is fed to a quadratic mixed Boolean programming (QBP) solver alongside an initial set of least-cost paths to be considered to service the required flows. The security layer receives the topology chosen by the functional layer and a security model to solve, with a mixed-integer programming (MIP) solver, the risk minimization problem. The output can be either infeasible (because its potential solution exceeds risk tolerances), in which case an analysis of the infeasibility produces a set of *equivalence classes* that is sent back to the functional layer to request the separation of specific flows that should not share paths. These equivalence classes generate additional constraints, known as Bender's cuts, that are added to the functional solver for a new iteration. The remainder of this section describes the major modules in Figure 1.

A. Functional Layer

The mathematical optimization model in the functional layer is a quadratic mixed binary programming model. In constraint programming the four main components are Inputs, Variables, Constraints, and an Objective function. Inputs are below.

Inputs

\mathcal{N} – the set of all network devices

\mathcal{E} – the set of edges (pair of vertices) connecting network devices

\mathcal{T} – the set of types of traffic to be routed

\mathcal{F} – the set of $(s, t, T) \in \mathcal{N} \times \mathcal{N} \times \mathcal{T}$ tuples defining desired

traffic flows of type T from source node s to sink node t .

$D(f) : \mathcal{F} \rightarrow \mathbb{R}$ – the actual demand for each flow $f \in \mathcal{F}$

$\mathcal{C} \subseteq 2^{\mathcal{N}}$ – a subset of sets of network devices

$\mathcal{R} \subseteq \mathcal{C} \times \mathcal{C}$ – pairs (c_1, c_2) of equivalence classes that segregate traffic from c_1 to c_2 .

\mathcal{P} – the set of all paths

$P(e) : \mathcal{E} \rightarrow \mathcal{P}$ – the set of all paths containing edge e

$P(n) : \mathcal{N} \rightarrow \mathcal{P}$ – the set of all paths containing node n

$P(c) : \mathcal{C} \rightarrow \mathcal{P}$ – the set of all paths containing a node in c

$N(p) : \mathcal{P} \rightarrow \mathcal{C}$ – the set of nodes appearing in path p

$P(s, t) : \mathcal{N} \times \mathcal{N} \rightarrow \mathcal{P}$ – the set of all paths $s \rightarrow t$

$cap(e) : \mathcal{E} \rightarrow \mathbb{R}$ – gives the capacity of an edge e .

Variables

$active_{p,T} \in \{0, 1\}$, – for every path $p \in \mathcal{P}$ and traffic type $T \in \mathcal{T}$, indicates whether path p carries traffic of type T

$flow_{p,T} \in \mathbb{R}_{\geq 0}$ – for every path $p \in \mathcal{P}$ and traffic type $T \in \mathcal{T}$, amount of flow of type T that is sent along path p

$equiv_{c,n} \in \{0, 1\}$ – does node $n \in \mathcal{N}$ appear in an active path together with a node in equivalence class c

$share_{c_1,c_2,n} \in \{0, 1\}$ – indicates whether node $n \in \mathcal{N}$ appears on any active path with nodes in classes $c_1, c_2 \in \mathcal{C}$. Namely,

$$share_{c_1,c_2,n} \Leftrightarrow n \in \left(\begin{array}{c} \left(\bigcup_{p \in P(c_1): active_{p,*}} N(p) \right) \\ \cap \\ \left(\bigcup_{p \in P(c_2): active_{p,*}} N(p) \right) \end{array} \right)$$

$active_{p,*} = 1$ if there is a type $T \in \mathcal{T}$ where $active_{p,T} = 1$

$load_n \in \mathbb{R}$ – the amount of flow that goes through node n

$loadObj$ – the sum of squares of all $load_n$ variables.

Constraints

$$\sum_{p \in P(e), T \in \mathcal{T}} flow_{p,T} \leq cap(e), \quad \forall e \in \mathcal{E} \quad (1)$$

$$\sum_{p \in P(s,t)} flow_{p,T} \geq D(s, t, T), \quad \forall (s, t, T) \in \mathcal{F} \quad (2)$$

$$active_{p,T} = 1 \rightarrow flow_{p,T} \geq 1, \quad \forall (s, t, T) \in \mathcal{F}, p \in P(s, t) \quad (3)$$

$$\sum_{p \in P(s,t)} active_{p,T} = 1, \quad \forall (s, t, T) \in \mathcal{F} \quad (4)$$

$$equiv_{c,n} = \bigvee_{p \in P(n) \cap P(c)} (active_{p,T}), \quad \forall T \in \mathcal{T}, n \in \mathcal{N}, c \in \mathcal{C} \quad (5)$$

$$share_{c_1,c_2,n} = equiv_{c_1,n} \wedge equiv_{c_2,n}, \quad \forall n \in \mathcal{N}, (c_1, c_2) \in \mathcal{R} \quad (6)$$

$$load_n = \sum_{p \in P(n), T \in \mathcal{T}} flow_{p,T}, \quad \forall n \in \mathcal{N} \quad (7)$$

Equation 1 enforces the edge capacity constraint to service the demand of all paths flowing through it. Equation 2 ensures that enough capacity is available to meet the demand of an (s, t, T) flow. Equation 3 ensures that some non-zero capacity is used if a specific path is activated (conversely, an inactive path can only have a 0 flow). Equation 4 states that a single path should be chosen to service a given flow $f \in \mathcal{F}$. Equations 5 define the auxiliary variables $equiv_{c,n}$ as true if and only if node $n \in \mathcal{N}$ appears on an active path sharing

a node with the equivalence class $c \in \mathcal{C}$. Equation 6 defines an active path that shares at least one node with two classes. Finally, equation 7 defines the load of a node as the sum of the flows associated to active paths passing through node n .

Objective

$$\min \left(\begin{array}{l} \alpha_0 \sum_{p,T} \text{len}(p) * \text{flow}_{p,T} \\ \alpha_1 \sum_{(c_1,c_2) \in \mathcal{R}, n \in \mathcal{N}} (\text{share}_{c_1,c_2,n} - 1) \\ \alpha_2 \sum_{n \in \mathcal{N}} (\text{load}_n)^2 \end{array} + \right) \quad (8)$$

The objective function 8 in this model is a weighted sum of three terms. The first term captures the total flows which are penalized by the length of the path used to dispatch those flows (such policies are codified in OSFP [51] and BGP practice [52]). The second term gives a unit credit each time equivalence classes on the segregation list \mathcal{R} do not share a node. (Due to this term, the objective value of the final solution may change between iterations of the functional layer.) The third and final term contribute to a bias towards solutions that achieve load balancing thanks to the quadratic component which heavily penalizes nodes with large loads.

Solving the Functional Model The functional model starts with empty sets \mathcal{C} and \mathcal{R} which are augmented with each iteration of the framework. New sets of nodes are added to \mathcal{C} and new segregation rules are added to \mathcal{R} (by the security layer). In the current implementation, least cost paths between pairs of nodes s, t are not generated “on demand”. Instead, the generation is limited to the first best k such paths, for increasing values of k . This process will ultimately be improved to use column generation techniques [16].

B. Risk Calculation

After the functional layer finds an optimal solution, it passes this solution to the risk calculation procedure. This input is the set of active paths. This module calculates the effective risk to the network for each path and traffic type.

Inputs

$\text{risk}(n, T) : \mathcal{N} \times \mathcal{T} \rightarrow \mathbb{R}$ – the risk inherit to network device n for traffic of type T ($\text{risk}(n, T) \geq 1$)

$d_k(n) : \mathcal{N} \rightarrow 2^{\mathcal{N}}$ – the set of nodes at a distance at most k from n in the logical topology

Calculation

Given an *active* path $p \in \mathcal{P}$ with source s and sink t , the calculation proceeds by partitioning the set of nodes of the path into three segments: the nodes “close” to the source s , “close” to the sink t and the nodes “in between”. Closeness is characterized by the function d_k and is meant to capture any connected node over the logical topology which sits no more than k hops away. Given this partition, $\text{flowRisk}(p, T)$ is:

$$\text{flowRisk}(p, T) = \sum_{i \in d_2(s) \cup d_2(t)} \text{risk}(i, T)^2 + \sum_{i \in N(p) \setminus (d_2(s) \cup d_2(t))} \text{risk}(i, T)^2$$

We use $k = 2$ to model nodes on the same LAN. The rationale is to impart to source s and sink t risk resulting from *lateral movement* of attacks. All other nodes contribute to the overall

path risk in proportion to the square of their own risks. We expect in most networks for $d_2(s)$ (and $d_2(t)$) to include nodes not directly on the path (like nodes on the same LAN). The input path risk calculation $\text{flowRisk}(p, T)$ is modular and can be augmented using other risk calculation methods.

C. Security Layer

The mathematical optimization model in the security layer is a mixed integer programming model. We similarly present the inputs, variables, constraints, and objective for the security layer. Its inputs are given below. Also note that all the variables from the functional model are *constants*.

Inputs

$\text{mem}(n) : \mathcal{N} \rightarrow \mathbb{R}$ – the memory resources of SDN device n
 $\text{fwCost}(T) : \mathcal{T} \rightarrow \mathbb{R}$ – the memory footprint for a firewall blocking traffic type T

piCost – the memory footprint for a packet inspection post
 fwComp – the complexity footprint for adding a firewall
 piComp – the complexity footprint for adding a packet inspection post to the network

$\text{penalty}(p, T) : \mathcal{P} \times \mathcal{T} \rightarrow \mathbb{R}$ – the penalty for blocking a unit of flow of type T along path p

$\text{rank}(n, p) : \mathcal{N} \times \mathcal{P} \rightarrow \mathbb{Z}$ – the position of node n in path p

$\text{flowRisk}(p, T) : \mathcal{P} \times \mathcal{T} \rightarrow \mathbb{R}_{\geq 0}$ – above risk calculation

Variables

$\text{fw}_{n,T} \in \{0, 1\}$ – does a firewall block traffic type T at n

$\text{pi}_n \in \{0, 1\}$ – is there packet inspection at network device n

$\text{fwOR}_{n,T} \in \{0, 1\}$ – is there a *block everything* or *block traffic of type T* firewall at network device n

$\text{fwOP}_{p,T} \in \{0, 1\}$ – is there a firewall on path p

$\text{piOP}_p \in \{0, 1\}$ – is there packet inspection on path p

$\text{rf}_{p,T} \in [0, 1]$ – risk factor for path $p \in P(s, t)$ servicing flow $(s, t, T) \in \mathcal{F}$

$\text{RMfw}_{p,n,T} \in [0, 1]$ – used in the riskFactor calculation

$\text{RMpi}_{p,n,T} \in [0, 1]$ – used in the riskFactor calculation

Constraints

$$\text{fwOR}_{n,T} = \text{fw}_{n,T} \vee \text{fw}_{n,*}, \forall n \in \mathcal{N}, T \in \mathcal{T} \quad (9)$$

$$\sum_{T \in \mathcal{T} \cup \{*\}} \text{fwCost}_T \cdot \text{fw}_{n,T} + \text{piCost} \cdot \text{pi}_n \leq \text{mem}_n, \forall n \in \mathcal{N} \quad (10)$$

$$\text{piOP}_p = \bigvee_{n \in N(p)} (\text{pi}_n), \forall T \in \mathcal{T}, p \in \mathcal{P} : \text{active}_{p,T} \quad (11)$$

$$\text{fwOP}_{p,T} = \bigvee_{n \in N(p)} (\text{fwOR}_{n,T}), \forall T \in \mathcal{T}, p \in \mathcal{P} : \text{active}_{p,T} \quad (12)$$

$$\text{fwOP}_{p,T} + \text{piOP}_p \geq 1, \forall T \in \mathcal{T}, p \in \mathcal{P} : \text{active}_{p,T} \quad (13)$$

$$\text{RMfw}_{p,n,T} = 1 - (.5)^{\text{rank}(n,p)} \cdot \text{fwOR}_{n,T}, \quad (14)$$

$$\forall p \in P(s, t), n \in N(p), (s, t, T) \in \mathcal{F}$$

$$\text{RMpi}_{p,n,T} = 1 - 0.1 \cdot (.5)^{\text{rank}(n,p)} \cdot \text{pi}_n, \quad (15)$$

$$\forall p \in P(s, t), n \in N(p), (s, t, T) \in \mathcal{F}$$

$$r_{f_p,T} = \min \bigcup_{n \in N(p)} \{RMfw_{p,n,T}, Rmpi_{p,n,T}\}, \quad (16)$$

$$\forall T \in \mathcal{T}, p \in \mathcal{P} : active_{p,T}$$

Equation 9 is used to define the presence of any kind of firewall at a node n . Equation 10 ensures that the memory footprint in SDN node n for the deployment of the firewall and the packet inspection logic does not exceed the device memory. Equation 11 links the packet inspection on a path with the presence of packet inspection on any node along the active path. Equation 12 plays a similar role for the presence of a firewall on a path. Equation 13 ensures that a firewall or a packet inspection is deployed somewhere on an active path. Equation 14 defines the minimum risk factor associated to a firewall. The earlier on the path the firewall is deployed, the lower the risk. Equation 15 similarly defines the minimal risk. Equation 16 defines the composite risk factor.

Objective

$$\min \left(\begin{aligned} &\beta_0 \left(\sum_{n,T} fwComp \cdot fw_{n,T} + \sum_n piComp \cdot pi_n \right) + \\ &\beta_1 \sum_n load_n \cdot pi_n + \\ &\beta_2 \sum_{p,T} penalty(p,T) \cdot flow_{p,T} \cdot fwOP_{p,T} + \\ &\beta_3 \sum_{p,T} flowRisk_{p,T} \cdot r_{f_p,T} \end{aligned} \right) \quad (17)$$

The objective function defined in equation 17 is a weighted sum of four distinct terms that focus on minimizing the network complexity based on security resources deployed, the load induced by inspection posts, the amount of desirable flow dropped due to firewall placement and finally the residual risk. This model is a classic mixed integer programming formulation.

D. Result Analysis

The result analysis layer tries to generate cuts for the functional layer with the goal of improving both functionality and security. To generate cuts, this layer will form equivalence classes of network nodes and pass back certain pairs of these classes, one at a time, to the functional layer. Each pair of classes describes a segregation rule, or a cut, to which to functional layer will adhere to as much as possible.

After the functional and security layers are re-optimized using the most recent cut, the result analysis layer determines whether the cut was beneficial or harmful based on the objectives of each layer. If the cut is deemed to have been beneficial, we permanently keep it as a constraint, repopulate the cut queue, and continue the process.

If the cut is deemed to have been harmful, it is removed from the functional layer's constraint pool. Then the next cut in the queue will be passed back to the functional layer. If the cut queue is empty, the feedback mechanism terminates and we output the best solution found.

We note that since this process only provides pairs of nodes it is a heuristic, it may be necessary for many nodes to simultaneously be separated to arrive at a global optimum. This mechanism performed well in our experiments.




Fig. 3. Order 4 Fat-tree with 2 gateway switches at the top and 2 hosts per edge switch.

E. Layer Coordination

It is valuable to review how the layers coordinate. The functional layer sends to the security layer a set of paths that implements the routing within the network to serve the specified flows while satisfying a set of segregation requirements. The security layer first computes *risks* for these paths (in polynomial time) based on its knowledge of the traffic. The paths, their risk and the security model are then tasked with deploying packet inspection apparatus as well as firewalls within that logical topology to monitor the traffic and block threats (risky traffic). Once the security model is solved to optimality, an analysis can determine whether the proposed logical topology is beneficial or not (w.r.t. its objective) and even suggest further equivalence classes for network nodes as well as segregation rules to be sent back to the functional layer for another iteration. Fundamentally, the coordination signal boils down to additional equivalence classes to group nodes together with segregation rules to separate paths that include network nodes in “antagonistic” equivalence classes.

F. Outputs

When the set of potential cuts is empty, the proposed configuration can be parsed and translated into SDN language fragments to be deployed on the network devices in order to obtain the desired logical network topology put forth by our framework.

IV. EXPERIMENTAL SETUP

A fundamental component of our work is the separation of the physical and logical networks. Our framework has potential in applications where many different logical topologies are possible from a single physical topology. Physical topology is an input to our framework and the empirical evaluation is based on a popular topology: Fat-Tree [53].

A Fat-Tree (mid 1980's) is shown in Figure 3. The instance of Fat-Tree we use was introduced by [53]. The updated design avoids bottlenecks through multiple equal capacity links between layers. This design uses four layers of switches: gateway, core, aggregate and edge with hosts connected at the bottom.

A. DDoS Scenario

Here we discuss the network setup simulating an emergent vulnerability/active attack. We select two hosts that are being targeted on a specific type of traffic, resulting in a significant increase in their risk for that type of traffic. Within our sample networks, we consider having two main types of devices: switches/routers and hosts. In order to model traffic between internal and external entities we utilize two gateway switches which represent the boundary of our network. For generality we consider two traffic types A and B which could represent any types of traffic such as web and storage. We also classify traffic as internal and external, with external traffic traversing one of the gateways. We allow only half of our hosts to communicate with external sources by allowing them to connect to one of the two gateways. Additionally all hosts are involved in internal communications. In both networks we have 16 hosts, we generated 60 flows with 44 of them being internal and 16 being external. We run multiple experiments providing the framework increasing numbers of starting paths between source and destination (from the priming procedure) to determine the impact on the solution quality.

Our implementation was built in Python 3.6 using the Gurobi 8.1 optimization library [54]. The experiments were run on a machine running Ubuntu 18.04 and equipped with an Intel Core i9-8950HK processor operating at 2.90 GHz with a 12 MB Cache and 32 GB of physical memory.

V. EVALUATION/RESULTS

We begin with the model’s resolution on a 16 host Fat-Tree using the 10 shortest paths per source and destination pair to prime the optimization. We assume equal demand for each of the 60 flows and two high risk hosts (in red in Figure 4). In the first iteration, the functional layer of the framework generates a candidate topology and passes this solution to the security layer. The security layer then calculates the network risk and deploys firewalls (represented with rectangles).

In the first iteration, the gateway on the right serves both low risk (shown in green) and high risk servers (shown in red). Deploying a firewall on this right gateway significantly reduces the network risk and is the output of the security layer. Importantly, this results in collateral damage as the low risk flows in green are also blocked. The first iteration through the framework deploys 3 firewalls to block 12 flows, 8 of which are high risk.

Iteratively, the security layer proposes separation of these collateral nodes from the high-risk nodes. The solution of the last framework iteration (493 candidates cuts are proposed, 40 prove beneficial) is shown on the right. This configuration routes all high risk flows through one core switch where a firewall is now deployed. Meanwhile, all the low risk flows access the gateways through a separate core switch.

Complete experimental results are shown in Table I. Our experiment consisted of running the DOCSDN framework the same physical topology with 60 flows being randomly selected for each experiment. We ran six experiments corresponding to a different bounded number of paths ($\{10, 20, 30, 40, 50, 100\}$)

being used to prime the functional layer for each source-destination pair. Ultimately, this process will be dynamic and use column-generation. A few observations are in order:

- The overall objective reflects both the functional and security layers. The other objective rows refer to each layer objective individually. The network risk rows quantify the risks and their change as the optimization proceeds. For instance, for the 10 paths benchmark, the risk degrades from 10425 to 11646 or 11.7% as a result of supporting an additional 4 good flows.
- The “nodes explored” rows indicate of the size of the branch and bound tree and remains quite modest throughout.
- All runs blocked all flows that contain a high risk server while preserving the low risk flows. All experiments delivered final configurations that preserved the same low-risk flows. We therefore hypothesize that a column-generation would quickly settle down and prove that no additional path can improve the quality of the solution. It is nonetheless interesting that adding more paths does not negatively impact the overall runtime.
- The objective functions of the functional and security layers use “scores” meant to ease the interplay between the two. Yet, it is wise to consult the *raw* properties of the solutions to appreciate the impact of the optimization. In particular, the number of flows blocked and the network risks. What is readily apparent is that improving functionality induces a slight degradation in the network risks, underlying the conflicting nature of the two objectives. The individual objective scores while moving in the correct direction are not to be viewed as stand alone metrics to determine solution quality but rather inter layer communications indicating improvement or decline from a functional or security perspective.
- The variance in time, iterations and number of cuts produced by each experiment is due to symmetries in the formulation. Solutions that are symmetric in the functional layer may not be symmetric in the security layer and induce slightly different solutions there. This is especially true for a Fat-tree network due to it’s built in redundancy/symmetry.
- Beneficial cuts reflects the number of segregation proposals from the security layers that are adopted by the functional layer (these cuts remove the current best feasible solution). harmful cuts are segregation proposals that do not “cut” the current best feasible solution or worsen the functional solution.

VI. CONCLUSION

Our framework is portable with respect to network risk assessment. Since the risk calculation/analysis is decoupled from the optimization model, the framework can be combined with any procedure that calculates risk on a per path basis. Along with this procedure, the other requirements for implementing a different risk mechanism:




Fig. 4. Illustration of the Fat-Tree network after the first pass through both layers of the framework (left) and the final solution (right). Note: Firewalls are depicted with rectangles, red nodes represent high risk nodes, green nodes represent nodes that are initially blocked and recovered in the final solution, utilizing the updated routes shown in dashed lines. The modification of the logical topology allows for more intelligent firewall placement balancing both functionality and security.

	10 paths	20 paths	30 paths	40 paths	50 paths	100 paths
Initial Flows Blocked	12	12	12	12	12	12
Final Flows Blocked	8	8	8	8	8	8
Initial Overall Objective	7874	7868	7870	7868	7854	7869
Final Overall Objective	3460	5555	4080	5235	4915	5761
Initial Functional Objective	2012	2012	2012	2012	2012	2012
Final Functional Objective	-6436	-2246	-5196	-2885	-3526	-1825
Initial Security Objective	13735	13724	13729	13723	13696	13726
Final Security Objective	13356	13356	13356	13356	13356	13348
Initial Network Risk	10425	10414	10419	10413	10386	10416
Final Network Risk	11646	11646	11646	11646	11646	11638
Functional Nodes Explored	370	55	206	83	510	46
Security Nodes Explored	1922	1650	152	30	28	19
Beneficial Cuts	40	20	34	23	26	18
Harmful Cuts	453	70	237	81	68	74
Iterations Needed	494	91	272	105	95	93
Time in Model (s)	283	40	319	112	76	273

TABLE I

EXPERIMENTAL RESULTS FROM APPLYING THE DOCSDN FRAMEWORK TO AN ORDER 4 FAT-TREE. EACH COLUMN REFERS TO A SEPARATE EXPERIMENT WHERE THE NUMBER OF PATHS PER SOURCE-DESTINATION PAIR GIVEN TO THE FRAMEWORK WERE VARIED.

- A way of evaluating how risk changes due to the deployment of network defenses
- The ability to propose candidate cuts that can be passed to the functional layer.

Our results show it is possible to effectively, automatically, and quickly find a network configuration that meets multiple conflicting properties. Our framework is modular, enabling integration of new desired properties. DOCSDN will allow network administrators to effectively prioritize and choose their desired properties. The efficiency of DOCSDN is enabled by the feedback/interplay between the functional and security optimization layers.

REFERENCES

- [1] A. Khurshid, W. Zhou, M. Caesar, and P. Godfrey, "Veriflow: Verifying network-wide invariants in real time," in *Proceedings of the first workshop on Hot topics in software defined networks*. ACM, 2012, pp. 49–54.
- [2] P. Neves, R. Calé, M. R. Costa, C. Parada, B. Parreira, J. Alcaraz-Calero, Q. Wang, J. Nightingale, E. Chirivella-Perez, W. Jiang *et al.*, "The SELFNET approach for autonomic management in an NFV/SDN networking paradigm," *International Journal of Distributed Sensor Networks*, vol. 12, no. 2, p. 2897479, 2016.
- [3] B. Schneier, "Attack trees," Blog, 1999. [Online]. Available: <http://tntlandforms.us/cs594-cns96/attacktrees.pdf>
- [4] R. Yu, G. Xue, V. T. Kilar, and X. Zhang, "Deploying robust security in internet of things," in *IEEE Conference on Computer and Network Security*, 2018.
- [5] R. Wang, D. Butnariu, J. Rexford *et al.*, "Openflow-based server load balancing gone wild," *Hot-ICE*, vol. 11, pp. 12–12, 2011.
- [6] R. Skowrya, A. Lapets, A. Bestavros, and A. Kfoury, "A verification platform for SDN-enabled applications," in *Cloud Engineering (IC2E), 2014 IEEE International Conference on*. IEEE, 2014, pp. 337–342.
- [7] R. Bixby, M. Fenelon, Z. Gu, E. Rothberg, and R. Wunderling, *System Modelling and Optimization: Methods, Theory, and Applications*. Kluwer Academic Publishers, 2000, ch. MIP: Theory and practice – closing the gap, pp. 19–49.
- [8] B. Fourer, "Amazing solver speedups," online, 2015, <http://bob4er.blogspot.com/2015/05/amazing-solver-speedups.html>.
- [9] J. F. Benders, "Partitioning procedures for solving mixed-variables programming problems," *Numerische Mathematik*, vol. 4, no. 1, pp. 238–252, 1962.
- [10] J. Hooker, "Logic-based Benders decomposition," *Mathematical Programming*, vol. 96, p. 2003, 1995.
- [11] —, *Logic-Based Methods for Optimization: Combining Optimization and Constraint Satisfaction*. John Wiley and Sons, 2000.
- [12] G. Codato and M. Fischetti, "Combinatorial Benders' cuts for mixed-integer linear programming," *Operations Research*, vol. 54, no. 4, pp. 756–766, 2006. [Online]. Available: <http://dx.doi.org/10.1287/opre.1060.0286>

- [13] C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. Savelsbergh, and P. H. Vance, "Branch-and-price: Column generation for solving huge integer programs," *Operations research*, vol. 46, no. 3, pp. 316–329, 1998.
- [14] H. Hijazi, T. W. K. Mak, and P. Van Hentenryck, "Power system restoration with transient stability," in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, ser. AAAI'15. AAAI Press, 2015, pp. 658–664. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2887007.2887099>
- [15] E. Lam and P. V. Hentenryck, "A branch-and-price-and-check model for the vehicle routing problem with location congestion," *Constraints*, vol. 21, no. 3, pp. 394–412, Jul. 2016. [Online]. Available: <http://dx.doi.org/10.1007/s10601-016-9241-2>
- [16] G. B. Dantzig and P. Wolfe, "Decomposition principle for linear programs," *Oper. Res.*, vol. 8, no. 1, pp. 101–111, Feb. 1960. [Online]. Available: <http://dx.doi.org/10.1287/opre.8.1.101>
- [17] P. Shaw, "Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems," in *Proceedings of Fourth International Conference on the Principles and Practice of Constraint Programming (CP'98)*. Springer Verlag, October 1998, pp. 417–431.
- [18] S. Kottler. (2018, March) February 28th ddos incident report. [Online]. Available: <https://githubengineering.com/ddos-incident-report/>
- [19] J. Ioannidis and S. M. Bellovin, "Pushback: Router-based defense against DDoS attacks," 2001.
- [20] —, "Implementing pushback: Router-based defense against DDoS attacks," in *NDSS*, vol. 2, 2002.
- [21] J. Mirkovic and P. Reiher, "A taxonomy of DDoS attack and DDoS defense mechanisms," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, pp. 39–53, 2004.
- [22] T. Peng, C. Leckie, and K. Ramamohanarao, "Survey of network-based defense mechanisms countering the DoS and DDoS problems," *ACM Computing Surveys (CSUR)*, vol. 39, no. 1, p. 3, 2007.
- [23] S. T. Zargar, J. Joshi, and D. Tipper, "A survey of defense mechanisms against distributed denial of service (DDoS) flooding attacks," *IEEE communications surveys & tutorials*, vol. 15, no. 4, pp. 2046–2069, 2013.
- [24] S. K. Fayaz, Y. Tobioka, V. Sekar, and M. Bailey, "Bohatei: Flexible and elastic DDoS defense," in *USENIX Security Symposium*, 2015, pp. 817–832.
- [25] B. Marczak, N. Weaver, J. Dalek, R. Ensafi, D. Fifield, S. McKune, A. Rey, J. Scott-Railton, R. Deibert, and V. Paxson, "Chinas great cannon," *Citizen Lab*, vol. 10, 2015.
- [26] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 267–280. [Online]. Available: <http://doi.acm.org/10.1145/1879141.1879175>
- [27] D. Kreutz, F. M. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmoly, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015.
- [28] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: enabling innovation in campus networks," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, 2008.
- [29] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story, and D. Walker, "Frenetic: A network programming language," *ACM Sigplan Notices*, vol. 46, no. 9, pp. 279–291, 2011.
- [30] H. Kim, J. Reich, A. Gupta, M. Shahbaz, N. Feamster, and R. J. Clark, "Kinetic: Verifiable dynamic network control," in *NSDI*, 2015, pp. 59–72.
- [31] J. Reich, C. Monsanto, N. Foster, J. Rexford, and D. Walker, "Modular SDN programming with Pyretic," *Technical Reprint of USENIX*, 2013.
- [32] R. Beckett, R. Mahajan, T. Millstein, J. Padhye, and D. Walker, "Network configuration synthesis with abstract topologies," in *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, 2017, pp. 437–451.
- [33] M. Corp, "Common vulnerabilities and exposures," December 2018. [Online]. Available: <https://cve.mitre.org>
- [34] NIST, "National vulnerability database," December 2018. [Online]. Available: <https://nvd.nist.gov>
- [35] G. Stoneburner, A. Y. Goguen, and A. Feringa, "SP 800-30. risk management guide for information technology systems," 2002.
- [36] W. Jansen, *Directions in security metrics research*. Diane Publishing, 2010.
- [37] S. Stolfo, S. M. Bellovin, and D. Evans, "Measuring security," *IEEE Security & Privacy*, vol. 9, no. 3, pp. 60–65, 2011.
- [38] R. Lippmann, J. Riordan, T. Yu, and K. Watson, "Continuous security metrics for prevalent network threats: introduction and first four metrics," Massachusetts Inst of Tech Lexington Lincoln Lab, Tech. Rep., 2012.
- [39] Y. Cherdantseva, P. Burnap, A. Blyth, P. Eden, K. Jones, H. Soulsby, and K. Stoddart, "A review of cyber security risk assessment methods for scada systems," *Computers & security*, vol. 56, pp. 1–27, 2016.
- [40] R. P. Lippmann and J. F. Riordan, "Threat-based risk assessment for enterprise networks," *Lincoln Laboratory Journal*, vol. 22, no. 1, pp. 33–45, 2016.
- [41] K. Ingols, R. Lippmann, and K. Piwowarski, "Practical attack graph generation for network defense," in *Computer Security Applications Conference, 2006. ACSAC'06. 22nd Annual*. IEEE, 2006, pp. 121–130.
- [42] K. Kaynar, "A taxonomy for attack graph generation and usage in network security," *Journal of Information Security and Applications*, vol. 29, pp. 27–56, 2016.
- [43] S. Zhang and S. Malik, "SAT based verification of network data planes," in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2013, pp. 496–505.
- [44] R. Beckett, A. Gupta, R. Mahajan, and D. Walker, "A general approach to network configuration verification," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. ACM, 2017, pp. 155–168.
- [45] F. Rossi, P. v. Beek, and T. Walsh, *Handbook of Constraint Programming (Foundations of Artificial Intelligence)*. New York, NY, USA: Elsevier Science Inc., 2006.
- [46] P. Baptiste, C. Le Pape, and W. Nuijten, *Constraint-Based Scheduling*. Kluwer Academic Publishers, 2001.
- [47] H. Nagarajan, E. Yamangil, R. Bent, P. V. Hentenryck, and S. Backhaus, "Optimal resilient transmission grid design," in *PSCC*. IEEE, 2016, pp. 1–7.
- [48] G. Byeon, P. Van Hentenryck, R. Bent, and H. Nagarajan, "Communication-Constrained Expansion Planning for Resilient Distribution Systems," *ArXiv e-prints*, Jan. 2018.
- [49] T. Coatta and G. W. Neufeld, "Configuration management via constraint programming," in *CDS*. IEEE, 1992, pp. 90–101.
- [50] S. Layeghy, F. Pakzad, and M. Portmann, "SCOR: software-defined constrained optimal routing platform for SDN," *CoRR*, vol. abs/1607.03243, 2016. [Online]. Available: <http://arxiv.org/abs/1607.03243>
- [51] J. T. Moy, *OSPF: anatomy of an Internet routing protocol*. Addison-Wesley Professional, 1998.
- [52] P. Gill, M. Schapira, and S. Goldberg, "A survey of interdomain routing policies," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 1, pp. 28–34, 2013.
- [53] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, ser. SIGCOMM '08. New York, NY, USA: ACM, 2008, pp. 63–74. [Online]. Available: <http://doi.acm.org/10.1145/1402958.1402967>
- [54] G. Optimization, "Inc.,gurobi optimizer reference manual, 2015," *URL: http://www.gurobi.com*, 2014.