

Compiling Minimum Incremental Update for Modular SDN Languages

Xitao Wen*, Chunxiao Diao*, Xun Zhao†, Yan Chen*, Li Erran Li‡, Bo Yang§, Kai Bu§
*Northwestern University, †Tsinghua University, ‡Bell Labs, Alcatel-Lucent, §Zhejiang University
{xitaowen2015,chunxiaodiao2012}@u.northwestern.edu, ychen@northwestern.edu,
‡zhaoxun09@mails.tsinghua.edu.cn, †erranli@research.bell-labs.com,
§{ybo2013,kaibu}@zju.edu.cn

ABSTRACT

Measurement results show that updating rules on switches poses major latency overhead during the course of the policy update. However, current SDN policy compilers do not handle policy updates well and generate large amount of redundant rule updates, most of which modify only the priority field. Our analysis shows that the lack of knowledge on the rule dependency and the consecutively distributed priority numbers are the fundamental problems behind the redundancy. In this paper, we propose to tackle the problems through 1) an extended policy compiler that builds rule dependency along with the compilation, and 2) an online optimization algorithm that maintains a scattered priority distribution. Our preliminary evaluation demonstrates that our proposed patch can eliminate nearly all the priority updates.

Categories and Subject Descriptors

C.2.3 [Computer Communication Networks]: Network Operations—*Network management*

Keywords

Software-defined Networking; OpenFlow; Compiler; Incremental Update

1. INTRODUCTION

Control plane modules are dynamic. The forwarding policies generated by the controller modules often have to dynamically react to network events with changes to the forwarding behaviors. Static forwarding policies that are too large to fit in a single flow table may also need dynamic swapping in reaction to the changing traffic patterns [4]. Although SDN allows a centralized approach to modify the forwarding policies installed in distributed switches, it does not avoid the significant latency overhead in altering the states of switches. According to the recent measurement

results [3], state-of-the-art OpenFlow (OF) switches can only process tens to a few hundreds of flow-mod instructions per second, which implies that a full refresh of a 5K-length flow table could lead to tens of seconds of volatile inconsistent window on the data plane. Although recent proposals on consistent updates can eliminate the potential erroneous forwarding behaviors, they also significantly increase the latency overhead by introducing multi-stage synchronization. Therefore, given a policy change, it is desirable to generate rule updates as compact as possible.

Policy composition of SDN policy languages makes the generation of compact updates challenging. In fact, parallel composition and sequential composition interleave the dependency relation among many rules and end up with one rule dependent on the policies from multiple modules. Our experiments show that a single-rule insert often results in modifying over half of the rules in a flow table with a straightforward update strategy on NetKAT [1]. The obscure dependency among rules forces the compiler to reassign a priority value to a rule even when its content aligns with an existing rule. The same problem applies to all current compilers that support policy composition. In our observation, we identify two major types of rule updates: content updates and priority updates. Content updates, which involve the modification of the predicates or actions of the rules, are mostly direct results of the changing policy content. On the other hand, priority updates, which only modify the rule priorities, are mostly caused by poorly distributed priority levels and unnecessary priority shifts. Surprisingly, the priority updates often dominate the size of the total updates, implying that the poorly handled priority is the major blame for the inflated update size. In this paper, we focus on eliminating unnecessary priority updates.

Our observation reveals two fundamental problems that prevent the compilers from generating good priority updates.

1. **Missing rule dependency clue.** Although OF rules are associated with sequential priority values, the dependency relationship among rules actually forms a directed acyclic graph (DAG) [4]. Unlike the DAG representation that characterizes the minimum set of dependencies, the sequential representation actually brings in plenty of non-existent dependencies and oftentimes makes avoidable priority updates necessary. Figure 1 shows an example where the dependency clue is essential to eliminate unnecessary priority updates.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotSDN'14, August 22, 2014, Chicago, IL, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2989-7/14/08 ...\$15.00.

<http://dx.doi.org/10.1145/2620728.2620733>.




Figure 1: Comparison of incremental update with vs. without dependency clue. In (b), due to the unawareness of the independency of the left and right rule chains, the update generator has to honor the priority levels set by the compiler and modify the priority of Rule 5, 6 and 7. While in (c), the optimal update requires only inserting Rule 8 in a new priority level.

2. Consecutive rule priority values. All state-of-the-art SDN policy compilers generate rules with consecutive priority values for simplicity. However, consecutive priority values prevent an update generator from inserting rules between two adjacent priority levels without affecting either of them. In fact, OF specification allocates a 16-bit integer for flow priority that allows 65536 different priority levels, whereas the total number of priority levels is usually much smaller than 65536, implying the possibility to scatter priority values.

In order to eliminate unnecessary priority updates, in this paper we propose 1) to guide update generation with rule dependency, and 2) to actively maintain the value gaps between priority levels.

The challenge of the first goal rests on how to obtain the dependency DAG. Intuitively, we can restore the dependency DAG directly from the output rules of the compiler with an existing algorithm [4]. However, we find such algorithm entails unacceptable $O(n^4)$ worst case computation complexity, where n is the size of the flow table. For this reason, we explore to build the dependency DAG incrementally along the compilation process, which significantly reduces the number of header space arithmetic operations and thus leads to more efficient algorithms. With the dependency DAG, we can generate a provable minimum-size update with regard to continuous priority levels (i.e., new levels can be arbitrarily inserted between two levels).

When mapping to discrete priority values, priority updates are sometimes necessary to make room for new levels. Intuitively, scattered priority values reduce the chance of future priority shifts. Abstractly, to maintain scattered priorities requires an online optimization strategy that makes proactive priority updates to minimize the estimation of future priority shifts for an undetermined policy update sequence. In this paper, we propose *k-factor strategy* that maintains the lengths of all gaps between $[\frac{1}{k}, k]$ times the average gap length. *K-factor strategy* achieves *k-factor* gap distribution with $O(1)$ amortized cost of priority shifts.

Our incremental update compilation framework is general to all SDN policy language compilers. To quantitatively estimate the benefit, we implement an initial prototype as an extension to NetKAT policy compiler, and evaluate it through benchmarks. As shown in Figure 2, our experiments demonstrate that the dependency information and a scat-




Figure 2: The update sizes of 100 successive policy updates on trace fw1. Two percent of the filters are replaced in each update.

tered priority distribution enable the compiler to eliminate nearly all the priority updates, and reduce the update size by an order of magnitude.

2. BACKGROUND

Related Work.

Several SDN policy languages (e.g., Frenetic [2], NetCore [6] Maple [11] and NetKAT [1]) have been proposed in recent years. Generally, a policy language compiler takes high-level policy descriptions and generates flow tables that fulfill the semantics of the policies. Except Maple, most compilers do not provide any support for incremental policy updates. In practice, they simply compile the new policies and replace the entire flow table of each switch. A straightforward improvement can be made by updating only the rules whose content or priority changes. Yet for the aforementioned problems, a considerable number of unnecessary rule updates still have to be conducted. On the other hand, although Maple does not support policy composition, it introduces tree-style abstraction to support incremental flow table updates. However, Maple compiler still makes a large amount of priority updates due to the consecutive priority values.

Another related topic is the incremental updates to TCAM. Many proposals have been made to reduce the update cost, which is measured by the number of entry moves in TCAM, in various context in terms of match method




Figure 3: The intersection of Rule A and B is shadowed by Rule C. Thus A and B are independent, i.e. the vertical order does not change their match space.

(longest-prefix match, range match, *etc.*) and number of matching fields (single-match or multiple-match) [7, 8, 10, 5]. Architecturally, our incremental update problem of SDN flow tables is orthogonal but complementary to that of TCAM in a sense that minimizing the updates to a flow table reduces the number of updates to TCAM, while incremental updating TCAM reduces the cost of each TCAM update.

Rule Dependency.

Intuitively, two rules are independent when their priority order does not change the header space classification. It is straightforward that two rules with disjoint predicates are independent. However, the opposite statement is not always true. For example, in Figure 3 the predicates of A and B intersect at a non-empty area, which is shadowed by the predicate of another rule C with higher priority. In this case, although $A \cap B \neq \emptyset$, they do not depend on each other in terms of representing the correct semantics. Therefore, we define two rules as *directly independent* if and only iff the predicates of two rules are disjoint or their intersection is entirely shadowed by other rules they both depend on. This is different from the dependency defined in several other works [11, 4], which do not consider the shadowing area.

Further, we define the *indirect dependency*, or *dependency* for short, by taking a transitive closure on the direct dependency relations. Obviously, the dependency relation induces a strict partial order in a flow table, and therefore forms a DAG of rules. The dependency DAG reveals the inherent relationship among rules in a sense that it represents the minimum set of the priority constraints in order to keep the flow space classification semantics.

3. SOLUTION OVERVIEW

Ideally, a minimum policy update only involves the rules whose content has changed. Having the existing flow table and the updated one, it is easy to find the rules whose content (except priority level) holds. However, the priority values of these content-invariant rules may have changed in the new table. To avoid changing the content-invariant rules, we have to shift their priorities in the new table to the old values. Meanwhile, such shifts may require further changes on the rules that are dependent on the content-invariant rules, or the rules those content-invariant rules depend on. Therefore, the minimum set of rule dependency of the new table is the key in generating of the minimum policy update. Such logic leads us to the design of the update generation framework as shown in Figure 4.

Preserving Dependency in Compilation.

In our framework, the policy compiler outputs the rule dependency DAG to describe the dependency among rules.

Since the current compilers typically maintain the dependency via priority value, it is feasible to restore the dependency DAG from the flow table through calculation. However, due to the high computation complexity of the restoration, we opt for the alternative approach to track the dependency during compilation (§4).

Comparing Updated Flow Table with Existing One.

To find the minimum rule update, the comparer identifies all the content-invariant rules by comparing the updated flow table with the existing one. For those content-invariant rules, we will try the best to reuse their current priority values unless there is a new priority level inserted into two priority levels that are currently mapped to two consecutive priority values. Other newly arrived rules are tagged on the dependency DAG to feed the prioritizer. At the end, those newly arrived rules are realized either by modifying retired rules or inserting new ones.

Assigning Priority Values to Priority Levels.

Initially, the priority levels should be evenly scattered on the priority value space. Upon each update, the prioritizer assigns priority values to priority levels based on their old values and the new dependency relation. The prioritizer maintains the distribution of the priority gaps with certain online strategy.

4. PRESERVING RULE DEPENDENCY

There are two ways to obtain rule dependency: to restore from flow table after compilation, or to gradually construct during the compilation, especially the module composition. In our analysis, we find the dependency construction algorithm proposed in CacheFlow [4] has time complexity of $O(n^3) \sim O(n^4)$ with regard to the length of the flow table¹. Meanwhile, we explore the latter approach and find a series of $O(n^2)$ algorithms to build the dependency DAG during compilation. In this section, we describe these algorithms on the abstractions of NetKAT local compiler. The algorithms can be easily adapted to other compilers that conduct similar policy composition.

4.1 SDN Compiler Background

The compilation of high-level SDN policy languages typically contains three major tasks, as depicted in Figure 5. First, the *syntax compiler* transforms a module’s policy from the high-level language to an intermediate representation (IR) of rule tables. Second, the *composition compiler* combines the potentially conflicting rule tables from different modules into a consistent rule table, according to the composition relationship among modules. Third, the *rule generator* translates the IR of rule tables into OpenFlow-compatible flow table.

NetKAT uses OpenFlow Normal Form (ONF), a subset of NetKAT language, as the intermediate representation towards OpenFlow flow table. An ONF flow table is comprised of abstract ONF rules, or atoms. Structurally, ONF forms a degenerate if-else binary tree representing the abstract atom sequence. Conceptually, the predicate of each abstract atom can be seen as a flow space filter with an

¹The algorithm description and the complexity analysis can be found in the technical report [12].




Figure 4: Overview of the solution.





Figure 5: SDN policy compilation and our proposed patch.



(b) Hierarchical dependency of ONF atoms.

outline match and possibly some holes representing the effect of the atoms it depends on, as depicted in Figure 6.

The compilation of NetKAT can be mapped to the compilation framework in Figure 4. NetKAT compiler first transforms the predicates and actions into if-else binary trees, corresponding to the syntax compilation step. Then, the compiler eliminates the composition operators through ONF tree transformation, corresponding to the composition compilation step. Finally, the rule generator transforms the abstract atoms into concrete flow rules, corresponding to the rule generation step.

As depicted in Figure 4, our goal is to explicitly build the dependency DAGs along each step of the the compilation. Particularly, during the compilation each ONF fragment should be associated with a dependency DAG, whose vertices represent ONF atoms and edges represent the dependency between atoms.

4.2 Initiating Dependency DAG in Syntax Compilation

After syntax compilation, NetKAT policies are essentially ONF fragments connected by composition operators. Particularly, at this stage each ONF fragment contains only one ONF atom. Therefore, the dependency DAGs associated with all of the ONF fragments are the same single-node graphs.

4.3 Preserving Dependency DAG in Composition

At this stage, NetKAT compiler recursively eliminates the composition operators and gradually combines the ONF fragments into a single ONF flow table. In each recursion, the compiler combines two ONF fragments connected by a composition operator into a single ONF fragment. Since each composition operator also combines the dependency DAGs of two ONF fragments, the goal of our extension is to combine and maintain the dependency DAG during NetKAT policy composition.

Parallel Composition.

Generally, the result of parallel composition contains three parts of policies, i.e., $S_1 = P_1 - P_2$, $S_2 = P_2 - P_1$ and $S_3 = P_1 \cap P_2$, as shown in Figure 7. Since dependency essentially can be seen as the holes in the match space, intuitively each of the three parts inherits a subset of the dependency (or holes) from the operands. The parallel composition of two ONF fragments is calculated by taking cross-product of predicates and actions of the two operands [1]. Thus, each rule in S_1 (or S_2) originates from a rule in P_1 (or P_2). Because the match of the new rule is a subset of that of the original rule, it inherits a subset of the dependency from the original rule. Similarly, the rules in S_3 may inherit dependency from both operands. Therefore, a superset of the result dependency DAG can be obtained by taking the cross-product of the operands' dependency DAGs, i.e., we have $R_{i,m}$ depends on $R_{j,n}$ in the result iff either R_i^{left} depends on R_j^{left} and $m = n$, or R_m^{right} depends on R_n^{right} and $i = j$.

Then, the edges of the result DAG have to be validated. We can simply take the intersection of the two rules on both ends of an edge. The dependency is still valid if and only if the intersection is not empty.

At last, one special treatment has to be made on the holes in S_3 . Since these holes are compiled to new rules, we need to add dependency between them and S_3 .

Sequential Composition.

The result of sequential composition only contains one part of rules, i.e., $P_1 \cap P_2'$ in Figure 8. In fact, the sequence of two ONF fragments P_1 and P_2 can be seen as the union of P_1 and P_2' , which is the projection of P_2 prior to the actions of P_1 . Since OF actions preserve the dependency relations, the result ONF fragment inherits a subset of the dependency from both operands. In other words, considering a result rule $R_{i,j}$ that originated from $R_i^{\text{left}} \in P_1$ and $R_j^{\text{right}} \in P_2$, $R_{i,j}$ may depend on every rule that originates from R_i^{left} and R_j^{right} . Like in parallel composition, all the dependency




Figure 7: Parallel composition's effect on dependency DAG.




Figure 8: Sequential composition's effect on dependency DAG.




Figure 9: Priority composition's effect on dependency DAG.

relations have to be validated by checking the overlapping of the two relevant rules.

Priority Composition.

To demonstrate the flexibility of the dependency DAG, we also devise a new binary composition operator, priority composition. The semantics of the priority composition is like setting different priority levels for two operands, as depicted in Figure 9. In other words, the second operand only takes effect on the flows that do not match the first operand. Generally, the priority union of two ONFs is calculated by taking a literal union of the operands' ONF rules with proper priority configuration.

The compilation of priority composition is made possible by the flexibility of dependency DAG. As the example in Figure 9, we just need to set the rules from P_1 to be inferior to all the rules from P_2 . Then, similar with other compositions, we validate the added dependency by checking the overlapping. In addition, the holes have to be treated specially.

4.4 Maintaining Dependency DAG in Rule Generation

In the rule generation stage, the rule generator splits an ONF atom into one base OF rule plus multiple shadow OF rules (i.e., the holes). The effect to the dependency DAG is treated in Algorithm 1.

5. MAINTAINING SCATTERED PRIORITY VALUES

Assigning priority values is an online strategy: upon each update, the prioritizer assigns the priority values to all new rules and possibly some old rules without the knowledge of the future updates. Intuitively, a more evenly scattered distribution of priority values reduces the chance of future priority updates with the cost of proactive priority updates. In this section, we describe the k -factor strategy, which actively maintains all gap lengths within the range of $[\frac{1}{k}, k]$ times the average gap length l_a , where k is a configurable parameter within the range $[1, +\infty)$. The k -factor strategy costs amortized $O(1)$ maintenance updates to achieve its gap distribution.

Without loss of generality, we assume the update is an insert or delete of a single rule. K -factor strategy works as follows. It first assigns all content-invariant rules with old priority values.

- If the update is a single-rule insert, the priority level of the new rule must be located between two existing priority levels on the priority DAG. Thus, the

input : Dependency DAG $G = (V, E)$, Mappings from DAG vertices to OF base rule and shadow rules $M : v \rightarrow \{b, S\}$

output: New dependency DAG $G' = (V', E')$

$G' = (\emptyset, \emptyset);$

foreach Mapping $u_i \rightarrow \{b, S\}$ **do**

$V' = V' + \{v_b, v_s^1, \dots, v_s^{|S|}\};$

foreach Edge $e = (u, u_i) \in E, u \neq u_i$ **do**

$E' = E' + (u, v_b);$

end

foreach Edge $e = (u_i, u) \in E, u \neq u_i$ **do**

foreach $i \in [1, |S|]$ **do**

if $v_s^i.match \cap u.match \neq \emptyset$ **then**

$E' = E' + (v_s^i, u);$

end

end

end

foreach $i \in [1, |S|]$ **do**

$E' = E' + (v_b, v_s^i);$

end

end

Algorithm 1: Maintaining Dependency DAG in Rule Generation

prioritizer allocates a new priority value between the two levels, and the gap is halved. If the lengths of the new gaps are less than the lower bound limit l_a/k , the prioritizer must shift the neighboring levels to meet the limit. Particularly, denoting the new level as m th level, the prioritizer determines the least number of neighboring levels n that satisfies $\lfloor (\sum_{i=m-n}^m l_i) / n \rfloor \geq l_a/k$ or $\lfloor (\sum_{i=m-n}^m l_i) / n \rfloor \geq l_a/k$. Finally, the prioritizer shifts the priority values of the n neighboring levels to equally partition the gaps.

- The process is similar for a rule delete except that the prioritizer now tests the upper bound limit kl_a instead of the lower bound limit l_a/k .

The parameter k balances the evenness of the gap distribution and the maintenance cost. At one end of the spectrum, when k equals 1, the prioritizer always maintains a uniform priority distribution, which costs more priority shifts for maintenance. At the other end, if k is large enough, the prioritizer only make priority shifts when no middle value is available for assignment.

6. EVALUATION

In this section, we evaluate how much size of the rule updates our proposed techniques can reduce.

| | Total #Rules | 2% | | | 4% | | | 6% | | |
|------|-----------------|--------|-------|---------|--------|-------|---------|--------|--------|---------|
| | | Naive | DAG | Optimal | Naive | DAG | Optimal | Naive | DAG | Optimal |
| fw1 | 898 | 592.40 | 52.03 | 37.06 | 589.39 | 92.10 | 66.03 | 493.81 | 113.25 | 79.23 |
| fw2 | 376 | 215.06 | 19.94 | 14.78 | 240.21 | 31.81 | 23.36 | 251.64 | 44.37 | 32.31 |
| acl1 | 121 | 56.50 | 3.19 | 2.61 | 71.58 | 7.18 | 5.46 | 80.37 | 11.71 | 8.77 |
| acl2 | 594 | 91.17 | 8.00 | 6.30 | 420.76 | 78.08 | 55.58 | 346.50 | 84.45 | 60.84 |
| ipc1 | 303 | 119.79 | 11.69 | 9.05 | 130.38 | 18.67 | 13.94 | 173.18 | 32.69 | 23.96 |
| ipc2 | 243 | 100.00 | 6.98 | 5.27 | 105.47 | 12.30 | 9.27 | 173.67 | 28.99 | 20.76 |

Table 1: Number of rules and mean size of rule updates of 100 successive policy updates. Correspondingly 2, 4 or 6 percent of the filters are replaced in each round of policy update.

Methodology.

We implement a prototype of the update generation framework based on NetKAT policy compiler. Our extension comprises two components: the DAG generator and the k -factor prioritizer. The DAG generator constructs the dependency graphs from prioritized OF rules with an algorithm adapted from CacheFlow [4]. The k -factor prioritizer compares the new rules with the existing ones and assigns priority value to the incremental rules with the k -factor strategy. We slightly modify NetKAT to allow the bit-wise IP prefix matches in our policy benchmark. We set parameter k to 2.

We build the base policies from filter sets generated by ClassBench [9]. We generate policy updates by randomly replacing a portion (2%, 4% or 6%) of filters from the previous filter set. For each configuration, we evaluate 100 successive rounds of policy updates. We compare the update size generated by our framework with 1) a naive strategy that updates the *diff* of flow tables, and 2) the optimal updates that contain no priority updates.

Results.

Table 1 shows the mean size of the 100 rounds of updates for each configuration. And Figure 2 shows all the size comparison of all 100 rounds of updates at the configuration [fw1, 2%], which represents a typical case in our experiments. We observe that the update size of the DAG cases is close to the optimal cases and is usually one order of magnitude smaller than that of the naive cases.

7. CONCLUSION AND FUTURE WORK

In this paper, we propose an update generation framework that utilizes the rule dependency information to minimize the number of rules to be modified in a policy update. We present the key techniques in the framework including how to obtain dependency information and how to maintain the scattered distribution of priority values. We further evaluate the benefit of our framework through benchmarks.

For future work, we would like to implement the incremental dependency construction on NetKAT policy compiler and evaluate the runtime performance gain. We will also explore other priority assignment strategies, especially the one that makes constant priority shifts in worst cases.

Acknowledgement

This material is based upon work supported in part by the NSF under grant No. CNS-1219116.

8. REFERENCES

- [1] ANDERSON, C. J., FOSTER, N., GUHA, A., JEANNIN, J.-B., KOZEN, D., SCHLESINGER, C., AND WALKER, D. NetKAT: Semantic Foundations for Networks. In *Proceedings of POPL '14* (2014), ACM, pp. 113–126.
- [2] FOSTER, N., HARRISON, R., FREEDMAN, M. J., MONSANTO, C., REXFORD, J., STORY, A., AND WALKER, D. Frenetic: A network programming language. In *ACM SIGPLAN* (2011), vol. 46.
- [3] HUANG, D. Y., YOCUM, K., AND SNOEREN, A. C. High-fidelity Switch Models for Software-defined Network Emulation. In *Proceedings of HotSDN '13* (New York, NY, USA, 2013), ACM, pp. 43–48.
- [4] KATTA, N., REXFORD, J., AND WALKER, D. Infinite cacheflow in software-defined networks. Tech. rep., TR-966-13, Department of Computer Science, Princeton University, 2013.
- [5] MISHRA, T., AND SAHNI, S. DUO-Dual TCAM Architecture for Routing Tables with Incremental Update. In *IEEE Symposium on Computers and Communications* (2010), Citeseer, pp. 503–508.
- [6] MONSANTO, C., FOSTER, N., HARRISON, R., AND WALKER, D. A compiler and run-time system for network programming languages. In *ACM SIGPLAN Notices* (2012), vol. 47, ACM, pp. 217–230.
- [7] SHAH, D., AND GUPTA, P. Fast updating algorithms for TCAMs. *IEEE Micro* 21, 1 (2001), 36–47.
- [8] SONG, H., AND TURNER, J. NXG05-2: Fast Filter Updates for Packet Classification using TCAM. In *Global Telecommunications Conference, 2006. GLOBECOM'06. IEEE* (2006), IEEE, pp. 1–5.
- [9] TAYLOR, D. E., AND TURNER, J. S. ClassBench: A Packet Classification Benchmark. *IEEE/ACM Trans. Netw.* 15, 3 (June 2007), 499–511.
- [10] VAN LUNTEREN, J., AND ENGBERSEN, T. Fast and scalable packet classification. *Selected Areas in Communications, IEEE Journal on* 21, 4 (2003), 560–571.
- [11] VOELLMY, A., WANG, J., YANG, Y. R., FORD, B., AND HUDAK, P. Maple: simplifying SDN programming using algorithmic policies. In *Proceedings of SIGCOMM '13* (2013), ACM, pp. 87–98.
- [12] WEN, X., LI, L. E., DIAO, C., ZHAO, X., AND CHEN, Y. Compiling Minimum Incremental Update for Modular SDN Languages. Tech. rep., NU-EECS-14-01, Department of EECS, Northwestern University, 2014.