

Automating SDN Composition: A Database Perspective

Anduo Wang* Jason Croft‡

* Temple University ‡University of Illinois at Urbana-Champaign

ABSTRACT

To keep up with the complexity of SDN management, it is generally agreed that modular development of the controller software plays a key role. However, forming a *correct* modular composition is still a challenging task. An operator needs to understand the module internals and to manually wire inter-module interactions that often depend on the underlying packets. In contrast, this poster presents a novel database approach towards *automatic, packet-agnostic composition* out of *black-box modules*.

1. PROBLEM STATEMENT

To keep up with the complexity of SDN management, it is generally agreed that modular development of the controller software plays a key role. Advanced modular SDN platforms [1] enable operators to program their network goals — traffic forwarding, fault tolerance, resource provisioning, stateful middleboxes, service chains, and more — as independently created control modules. However, forming a *correct* modular composition, where multiple control modules collectively realize a coherent network behavior, is still a challenging task. An operator often needs to understand a member module’s internals to determine the fine-grained inter-module interactions in all execution runs. Thus, while the SDN paradigm simplifies *how to realize* individual control modules, it relies on the operator to determine *what is the correct composition*.

While there is little hope in building a general-purpose procedure that stitches together arbitrary software modules, we believe it is feasible to automate composition for a large set of SDN modules. We build on the insight that an SDN control module, regardless of the disparate language constructs being adopted, conceptually follows a common pattern of *checking and repair*. A module continuously *checks* the network states for violations of some invariants; if a violation is detected, the module computes a new network state to *repair* the network and restore the invariants. Based on this assumption, the crux to a correct SDN composition is to prevent any module repair from reconfiguring the network into a state that violates other modules.

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 ACM 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.

SOSR '17, April 03-04, 2017, Santa Clara, CA, USA

© 2017 ACM. ISBN 978-1-4503-4947-5/17/04...\$15.00

DOI: <http://dx.doi.org/10.1145/3050220.3060612>

Correctness criteria. First, we establish what constitutes a *correct* composition, identifying criteria that guide us through the correct composition. We propose defining composition correctness as a semantic-preserving property: a composite is correct if it respects the semantics (invariants) of every constitute module. Given a set of modules M , each $m \in M$ maintains (checks and repairs) some invariant i_m about the network, a composition out of these modules is correct if it continues to enforce every invariants $\bigwedge_{m \in M} i_m$.

To see the strength of this definition, consider two simple modules p, q in NetKAT [1] syntax, where \cdot represents sequential composition, and $+$ denotes parallel composition. p drops (0) any packets whose destination matches x ($\text{dst}=x$), and q sets the destination of some packets (those matching some_predicate) to x ($\text{dst} \leftarrow x$):

$$p \triangleq (\text{dst}=x) \cdot 0$$

$$q \triangleq (\text{some_predicate}) \cdot (\text{dst} \leftarrow x)$$

There are three different ways to compose p and q , namely $p+q$, $p \cdot q$, and $q \cdot p$. Both $p+q$ and $p \cdot q$ will drop packets originally destined to x but redirect some other packets (those matching some_predicate) to x . Only $q \cdot p$ will drop all packets destined to x (including packets redirected to x by q). If we view the semantics (invariant) of p as removing any packets destined to x , only $q \cdot p$ respects p ’s semantics. $p+q$ and $p \cdot q$ — by leaving some packets destined to x unfiltered — both violate q ’s intention. Thus, the only correct composition is $q \cdot p$.




Figure 1: Example network

Reference composition. While our criteria enable us to decide whether or not a composition is correct, it does not automatically construct one. To see the subtlety in constructing a correct composition, consider the network in Fig. 1. The network consists of three boxes — firewall FW , switch A , and load balancer LB — that connect two external hosts H_1, H_2 to two back-end servers S_1, S_2 sharing a public address (S). Also suppose the network is controlled by three modules independently developed: routing (rt) establishes routes between the clients and the public face of the servers; firewall (fw) filters packets originated from certain clients (H_1); and load balancer (lb) distributes and forwards packets from/to the servers. A possible high-level specification of fw, lb is illustrated in the following (we omit rt):

$$\text{fw} \triangleq \neg(\text{src}=H_1.\text{dst}=S) + \neg(\text{src}=S.\text{dst}=H_1)$$

$$\text{lb} \triangleq \text{if}_{(e_c, \text{lb}_1, \text{lb}_2)} \text{ where}$$

$$\text{lb}_1 \triangleq \text{src}=H_1.\text{dst}=S.\text{dst} \leftarrow S_1.\text{pt} \leftarrow 2 + \dots$$

$$\text{lb}_2 \triangleq \text{src}=S_1.\text{dst}=H_1.\text{src} \leftarrow S.\text{pt} \leftarrow 1 + \dots$$

Note that lb , with an auxiliary filter ($e_c \triangleq_{sw=FW.pt=1\dots}$) denoting packets entering the network from the clients, defines two distinctive repairing actions based on the pattern of the packets. For packets from the clients (e_c), lb_1 maps the public server address (s) to a particular private address (s_1), otherwise it restores the source address from a particular server to s . A correct composition, a *master program* that modularly combines these modules, is the following:

$if_(e_c, fw.(lb_1+rt), lb_2.fw.rt)$

This master program instructs the packets from the clients to be processed by the firewall first, followed by routing and load balancer ($fw.(lb_1+rt)$). For packets from the servers, the master program schedules the load balancer before the firewall ($lb_2.fw.rt$) so that the public address (s) can be properly restored for the firewall. In this composition, the member module lb can *not* be treated as a black-box. Rather, the operator needs to understand the *internals* of lb and to explicitly wire its interactions with fw, rt . It is also called the “decompose and re-compose” problem, and additional requirements can quickly add complexity.

In contrast to this manual subtle composition, our objective is to *automatically* generate a *traffic-agnostic composition* out of (possibly finer-grained) *black-box modules*, like the following equivalent composition ($e \triangleq_{e_s+e_c}$ denoting all packets entering the network edges):

$e.lb_2.fw.(lb_1+rt)$

2. A DATABASE SOLUTION

Recall the two simple modules p, q (§1). The correct composition $q.p$ can be derived from the fact that the network repair pushed by q can change the network into a state that activates p — a state that violates the invariants checked by p . We introduce behavioral dependency, denoted by $x \rightarrow y$ (or $y \leftarrow x$), to capture the dependence of module x on y : when x ’s repair can inadvertently introduce violations to y , x requires further (repairing) service from y . If $x \rightarrow y$, we compose by $x.y$; otherwise, if neither $x \rightarrow y$ nor $y \rightarrow x$, we compose by $x+y$.

This composition strategy generalizes to a set of modules M as follows: First, we compute a dependency graph D that represents all dependencies. Next, we develop a stratification algorithm to convert D into a correct composition.




Figure 2: Stratification of dependency graph (left). Constructing the reference composition based on stratification number (right).

Constructing composition (stratification algorithm). Fig. 2 (left) depicts the dependency graph for the modules $\{e, lb_1, lb_2, fw, rt\}$ collectively driving the network in Fig. 1. In the dependency graph, each module x (node) is associated with a stratification number ($sn(x)$ in red). Intuitively, this number decides the (sequential) ordering in




Figure 3: Automatic dependency discovery by SAT/UNSAT reasoning. Invariants θ, δ in the diamond, repairing updates in the square.

the modular composition. Modules with equal stratification numbers, on the other hand, are independent and are composed in parallel. More precisely, the stratification number is given by:

$$sn(x) = \begin{cases} \max(\{sn(y) \mid \text{if } y \text{ is a parent of } x \text{ in } D\}) + 1 \\ 1 \text{ if } x \text{ is the root of } D \end{cases}$$

Fig. 2 (right) shows the composition constructed from the modules’ stratification numbers. In general, we have:

$$\forall u, v, x, y \in M, sn(u) + 1 = sn(v) + 1 = sn(x) = sn(y) \\ comp(M) = \dots (u + v + \dots) \cdot (x + y + \dots) \dots$$

Generating the dependency graph. We build D by determining the behavioral dependency relation between every pair of modules. To determine if $x \rightarrow y$ (or $y \rightarrow x$), we leverage automatic reasoning based on database irrelevant updates. The key idea is to represent all network states as database tables [3]: the dataplane as shared, stored tables, while the modules operate on derived tables (database views); and to reduce a module x ’s operations — the checking and repairing of invariants — to a database query (view x_v) and a database update (view update x_u). This allows us to apply database irrelevance reasoning [2] to generate the dependency relation: For $\forall x, y \in M$, x depends on y ($x \rightarrow y$) if x ’s updates can cause a violation to y (x_u is relevant to y_v), but y ’s updates can never affect x (y_u is irrelevant to x_v).

Fig. 3 illustrates the reasoning process for deciding $lb_2 \rightarrow fw$: lb_2 update (insert $+rm(_, S, H_1)$ to reachability table rm) can violate fw because the invariant θ is satisfiable (SAT); but fw will not inadvertently affect lb_2 because the conjunction of δ and the repairs is not satisfiable (UNSAT).

While our composition method is discussed in a concrete SDN platform (i.e., NetKAT), we believe our strategy will apply to a large family of SDN modules — the particular choice of language is irrelevant; the enabling assumption is an SDN control’s check-repair structure.

Acknowledgment This material is based upon work supported by the NSF Grant CNS 1657285.

3. REFERENCES

- ANDERSON, C. J., FOSTER, N., GUHA, A., JEANNIN, J.-B., KOZEN, D., SCHLESINGER, C., AND WALKER, D. Netkat: Semantic foundations for networks. In *POPL ’14*.
- BLAKELEY, J. A., COBURN, N., AND LARSON, P.-V. Updating derived relations: Detecting irrelevant and autonomously computable updates. *ACM Trans. Database Syst.* 14, 3 (Sept. 1989), 369–400.
- WANG, A., MEL, X., CROFT, J., CAESAR, M., AND GODFREY, B. Ravel: A database-defined network. In *SOSR ’16*.