



Association for
Computing Machinery

acm



sigcomm



SIGCOMM 2018

BUDAPEST

TRIDENT

TOWARD A UNIFIED SDN PROGRAMMING FRAMEWORK WITH AUTOMATIC UPDATES



Kai Gao¹ Taishi Nojima² Y. Richard Yang^{2,3}

August 23, 2018


¹Tsinghua University ²Yale University ³Tongji University



SOFTWARE-DEFINED NETWORKING




SOFTWARE-DEFINED NETWORKING




SDN simplifies network management with logically centralized network control.

NETWORK FUNCTIONS



NETWORK FUNCTIONS




Network functions provide L7 information by extracting “state”.

PUTTING THEM TOGETHER

Integrating the information extracted by network functions into SDN programming enables **adaptive, cross-layer** network control.


- adaptive: react dynamically to traffic
- cross-layer: control traffic based on L2-L7 information



PUTTING THEM TOGETHER

Integrating the information extracted by network functions into SDN programming enables **adaptive, cross-layer** network control.

- adaptive: react dynamically to traffic
- cross-layer: control traffic based on L2-L7 information



Unified SDN Programming

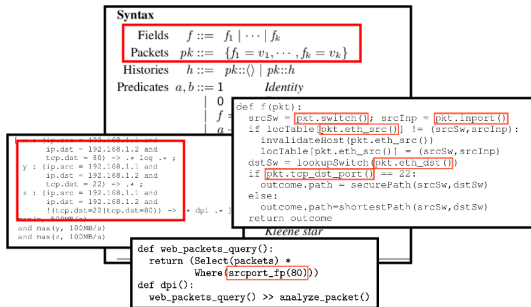
*What are the design challenges of
a unified SDN programming framework?*

&

Why are existing SDN programming frameworks not sufficient?

C1: INTEGRATING NETWORK FUNCTION STATE INTO SDN PROGRAMMING


- State-of-the-art SDN programming languages support L2-L4 programming naturally as all L2-L4 information is contained in every single packet.



Examples from NetKAT (Anderson et al.), Frenetic (Foster et al.), Maple (Voellmy et al.) and Merlin (Soulé et al.)

C1: INTEGRATING NETWORK FUNCTION STATE INTO SDN PROGRAMMING


- State-of-the-art SDN programming languages support L2-L4 programming naturally as all L2-L4 information is contained in every single packet.
- Network function states are L7 which are **NOT** contained in packet header fields. They can be **unknown** and **constantly updated by finite state machines**.



Examples from Kinetic (Kim et al.) and Resonance (Kim et al.)

C1: INTEGRATING NETWORK FUNCTION STATE INTO SDN PROGRAMMING

- State-of-the-art SDN programming languages support L2-L4 programming naturally as all L2-L4 information is contained in every single packet.
- Network function states are L7 which are **NOT** contained in packet header fields. They can be **unknown** and **constantly updated by finite state machines**.



Examples from Kinetic (Kim et al.) and Resonance (Kim et al.)


We need a simple abstraction to encode L7 information in SDN programming.

C2: CONSTRUCTING CONSISTENT, CORRELATED ROUTES

- Route constructions **may be required to be correlated**: routes cannot be calculated independently.

Requirement Case 1:

The return path be the inverse of the forward path (i.e., symmetry).




If the forward and return paths are computed independently using shortest path, the requirement will not be satisfied.

C2: CONSTRUCTING CONSISTENT, CORRELATED ROUTES

- Route constructions **may be required to be correlated**: routes cannot be calculated independently.

Requirement Case 1:

The return path be the inverse of the forward path (i.e., symmetry).



If the forward and return paths are computed independently using shortest path, the requirement will not be satisfied.


We need to systematically construct consistent, correlated routes.

HIGH-LEVEL PROGRAMMING ABSTRACTIONS IN TRIDENT



To address the aforementioned challenges

HIGH-LEVEL PROGRAMMING ABSTRACTIONS IN TRIDENT



To address the aforementioned challenges, Trident introduces

- **stream attribute**, to encode a network function state as if it is a header field so that programmers can select packets based on network function states,
- **route algebra**, a simple yet flexible abstraction to systematically construct consistent, correlated routes.

HIGH-LEVEL PROGRAMMING ABSTRACTIONS IN TRIDENT



Network function states are dynamic

- When the state of a finite state machine for a network function changes, the corresponding route should be updated to be consistent.
- Handling dynamicity manually is complex and error-prone.

HIGH-LEVEL PROGRAMMING ABSTRACTIONS IN TRIDENT




Network function states are dynamic

- When the state of a finite state machine for a network function changes, the corresponding route should be updated to be consistent.
- Handling dynamicity manually is complex and error-prone.


We need to automatically handle consistent updates!

HIGH-LEVEL PROGRAMMING ABSTRACTIONS IN TRIDENT



Trident introduces **live variable abstraction**

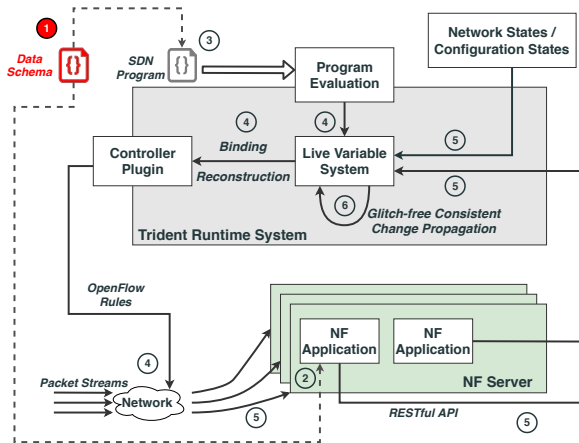
HIGH-LEVEL PROGRAMMING ABSTRACTIONS IN TRIDENT



Trident introduces **live variable abstraction** to handle dynamicity of both stream attributes and route algebra.

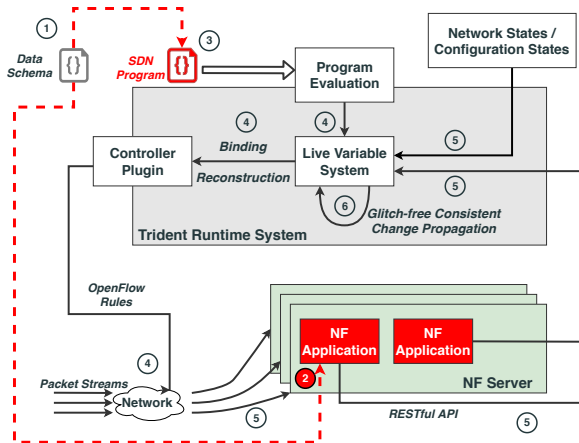
WORKFLOW OF TRIDENT

1 Network operator & programmer specifies data schema for network function states.



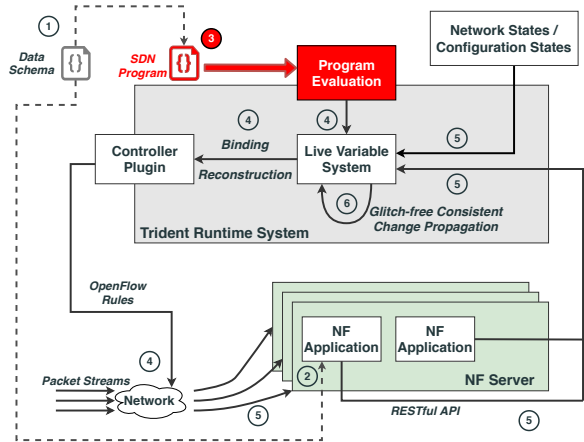
WORKFLOW OF TRIDENT

2 Network functions implement the schema.



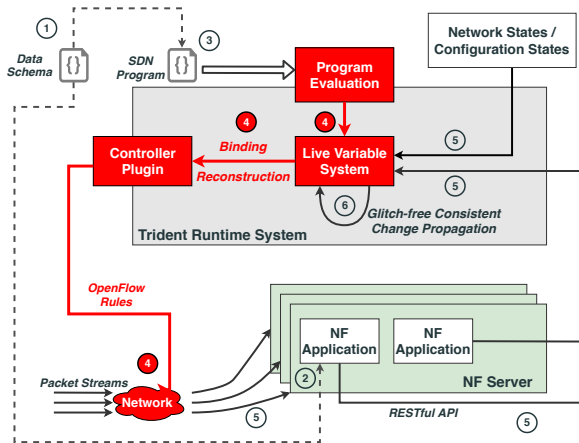
WORKFLOW OF TRIDENT

3 Network operator submits the program to Trident.



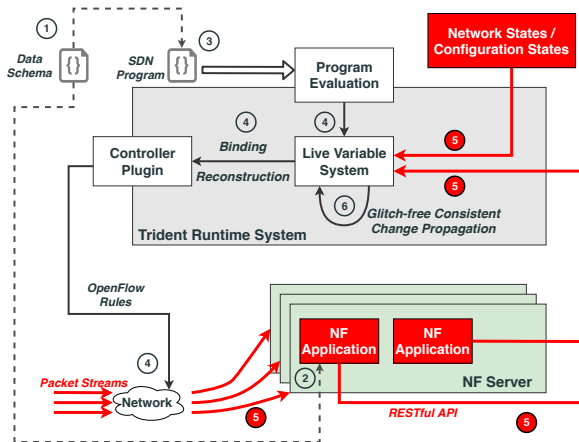
WORKFLOW OF TRIDENT

- 4 Trident evaluates the program and calculates the corresponding routes.



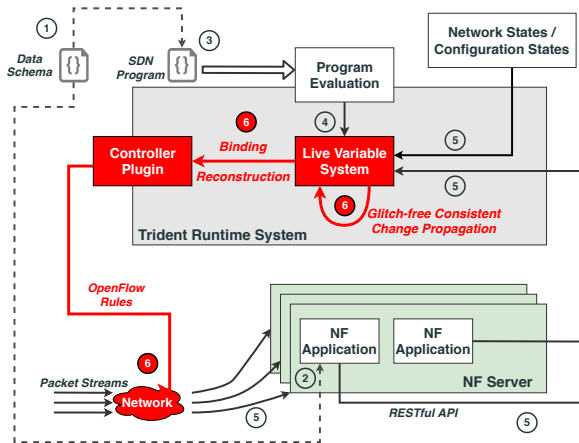
WORKFLOW OF TRIDENT

5 A change comes: a network function updates its state, a network state changes (e.g., a link fails), or a configuration is changed (e.g., a change to an access control list).



WORKFLOW OF TRIDENT

6 Trident automatically updates the routes for any change.



STREAM ATTRIBUTE: DETAIL

Observation: Different network function states are computed from different sets of packets.

Example

For example:

- HTTP URI: Computed from packets of the same TCP connection defined by TCP 5-tuple (e.g., **<10.0.0.2, 10.0.1.2, 1234, 80, tcp>**)
- Heavy hitter (source): Computed from packets with the same source IP address (e.g., **10.0.0.2**).

STREAM ATTRIBUTE: DETAIL

Define a stream attribute:

```
val http_uri = StreamAttribute[String]("HTTP_URI", TCP5TUPLE)
```

STREAM ATTRIBUTE: DETAIL

Define a stream attribute:

```
val http_uri = StreamAttribute[String] ("HTTP_URI", TCP5TUPLE)
```

- The type information, e.g., `String`, `Int`

STREAM ATTRIBUTE: DETAIL

Define a stream attribute:

```
val http_uri = StreamAttribute[String] "HTTP_URI", TCP5TUPLE)
```

- The type information, e.g., `String`, `Int`
- A descriptive name, e.g., `HTTP_URI`, `authenticated`

STREAM ATTRIBUTE: DETAIL

Define a stream attribute:

```
val http_uri = StreamAttribute[String] ("HTTP_URI", TCP5TUPLE)
```

- The type information, e.g., `String`, `Int`
- A descriptive name, e.g., `HTTP_URI`, `authenticated`
- The **stream type** (bit masks on packet header fields) specifying the set of packets to compute the network function state, e.g., `TCP5TUPLE`, `SRC_IPADDR`, `DST_IPADDR`

STREAM ATTRIBUTE: DETAIL

Use a stream attribute **just like a packet header**

1 `pkt.http_uri, pkt.authenticated, pkt.heavy_hitter, ...`

STREAM ATTRIBUTE: DETAIL

Use a stream attribute **just like a packet header**

1 `pkt.http_uri, pkt.authenticated, pkt.heavy_hitter, ...`

Stream attribute MAY have an **unknown** value.

Trident treats unknown values as valid and uses **Kleene's 3-valued logic** to select packets based on stream attribute.

Truth tables for \wedge and \vee in Kleene's 3-valued logic (T - True, F - False, U - Unknown)

\wedge	T	F	U
T	T	F	U
F	F	F	F
U	U	F	U

\vee	T	F	U
T	T	T	T
F	F	F	U
U	U	U	U

STREAM ATTRIBUTE: DETAIL

Use a stream attribute **just like a packet header**

```
1  pkt.http_uri, pkt.authenticated, pkt.heavy_hitter, ...
```

Stream attribute MAY have an **unknown** value.

Trident treats unknown values as valid and uses **Kleene's 3-valued logic** to select packets based on stream attribute.

```
1      // 3-way branch
2      if ((pkt.authenticated) && (pkt.http_uri === "www.xyz.com")) {
3          // true branch
4      } else {
5          // else branch
6      } unknown {
7          // unknown branch
8      }
```

```
1      // fallback branch
2      iff ((pkt.authenticated) && (pkt.http_uri === "www.xyz.com")) {
3          // true branch
4      } else {
5          // else and unknown branch
6      }
```

ROUTE ALGEBRA: DETAILS

Objective: Use well-structured, declarative expressions to specify the construction of consistent, correlated routes (motivated by prior studies such as **waypoint-based routing**¹ and **relational algebra**²).

- The basic unit of route algebra is **route set**.
- Each route set has a **network function indicator** to specify the symmetry requirements of network functions.

¹NetKAT (Anderson et al., POPL'14), Merlin (Soulé et al., CoNEXT'14), Propane (Beckett et al., SIGCOMM'16/PLDI'17), Genesis (Subramanian et al., POPL'17)

²EF Codd. "RELATIONAL COMPLETENESS OF DATA BASE SUBLANGUAGES". In: *Computer* (1972)

ROUTE ALGEBRA: DETAILS

Union (\cup)/Intersection (\cap)/Difference (\setminus)

Given two route set Δ_1 and Δ_2 , return the union/intersection/difference of Δ_1 and Δ_2 :

$$\begin{aligned}\Delta_1 \cup \Delta_2 &= \{r \mid r \in \Delta_1 \vee r \in \Delta_2\}, \\ \Delta_1 \cap \Delta_2 &= \{r \mid r \in \Delta_1 \wedge r \in \Delta_2\}, \\ \Delta_1 \setminus \Delta_2 &= \{r \mid r \in \Delta_1 \wedge r \notin \Delta_2\}.\end{aligned}$$

Union (\cup_{\sim})/Intersection (\cap_{\sim})/Difference (\setminus_{\sim}) by Equivalence

Given two route set Δ_1 and Δ_2 , return the union/intersection/difference of Δ_1 and Δ_2 using \in_{\sim} instead of \in :

$$\begin{aligned}\Delta_1 \cup_{\sim} \Delta_2 &= \{r \in \Delta_1 \cup \Delta_2 \mid r \in_{\sim} \Delta_1 \vee r \in_{\sim} \Delta_2\}, \\ \Delta_1 \cap_{\sim} \Delta_2 &= \{r \in \Delta_1 \cup \Delta_2 \mid r \in_{\sim} \Delta_1 \wedge r \in_{\sim} \Delta_2\}, \\ \Delta_1 \setminus_{\sim} \Delta_2 &= \{r \in \Delta_1 \cup \Delta_2 \mid r \in_{\sim} \Delta_1 \wedge r \notin_{\sim} \Delta_2\}.\end{aligned}$$

Concatenation (+)

Given two route sets Δ_1 and Δ_2 , return a new route set by concatenating all route pairs (r_1, r_2) in $\Delta_1 \times \Delta_2$ and removing the invalid ones:

$$\Delta_1 + \Delta_2 = \{r_1 + r_2 \mid r_1 \in \Delta_1, r_2 \in \Delta_2, dst_{r_1} = src_{r_2}\}.$$

Inversion (\asymp)

Given a route set Δ , return the inverse of $r \in \Delta$:

$$\asymp \Delta = \{\asymp r \mid r \in \Delta\}.$$

Preference (\triangleright)

Given two route sets Δ_1 and Δ_2 , return the *preferred* route. (If there is an equivalent route in Δ_1 , do not use the ones in Δ_2):

$$\Delta_1 \triangleright \Delta_2 = \{r \mid r \in \Delta_1 \vee (r \in \Delta_2 \wedge \nexists r' \in \Delta_1, r \sim r')\}.$$

Selection (σ)

Given a route set Δ and an evaluation function $f : R^* \mapsto \{0, 1\}$, return all routes in Δ that are evaluated as 1:

$$\sigma_f(\Delta) = \{r \in \Delta \mid f(r) = 1\}.$$

Optimal selection (\diamond)

Given one route set Δ and a routing cost function $d : R^* \mapsto \mathbb{R}$, return *any* route with the minimum value:

$$\diamond_d(\Delta) = \arg \min_{r \in \Delta} d(r).$$

Arbitrary selection ($*$)

Given one route set Δ , return a route set containing exactly one route r in Δ :

Please see the paper for detailed specification.

ROUTE ALGEBRA: EXAMPLE

- Route for a flow from a host to a gateway with link capacity preference (prefers 100 Gbps over 10 Gbps):

$$* \left(\sigma_{cap=100Gbps}(H : - : GW) \triangleright \sigma_{cap=10Gbps}(H : - : GW) \right)$$

Step 1: Compute the primary route set with high link capacity.

ROUTE ALGEBRA: EXAMPLE

- Route for a flow from a host to a gateway with link capacity preference (prefers 100 Gbps over 10 Gbps):

$$* \left(\sigma_{cap=100Gbps}(H : - : GW) \triangleright \sigma_{cap=10Gbps}(H : - : GW) \right)$$

Step 2: Compute the backup route set with low link capacity.

ROUTE ALGEBRA: EXAMPLE

- Route for a flow from a host to a gateway with link capacity preference (prefers 100 Gbps over 10 Gbps):

$$* \left(\sigma_{cap=100Gbps}(H : - : GW) \triangleright \sigma_{cap=10Gbps}(H : - : GW) \right)$$

Step 3: Combine them together with the **preference** operator.

ROUTE ALGEBRA: EXAMPLE

- Route for a flow from a host to a gateway with link capacity preference (prefers 100 Gbps over 10 Gbps):

$$* \left(\sigma_{cap=100Gbps}(H : - : GW) \triangleright \sigma_{cap=10Gbps}(H : - : GW) \right)$$

Step 4: Select only one route for unicast.

LIVE VARIABLE: DETAILS

Objective: Make dependency tracking and updates **transparent to programmers** (motivated by functional reactive programming³).

- Live variable is a **traceable data type** which stores the value and the computation process (i.e., dependencies and computation methods).
- Stream attribute and route algebra are just **higher abstractions** of live variables.
- The update of live variables satisfies the **glitch-free consistency**.

³Fran (Elliott and Hudak, IFIP'97), Dream (Margara and Salvaneschi, DEBS'14) and REScala (Drechsler et al., OOPSLA'14)

LIVE VARIABLE: DETAILS

Objective: Make dependency tracking and updates **transparent to programmers** (motivated by functional reactive programming³).

- Live variable is a **traceable data type** which stores the value and the computation process (i.e., dependencies and computation methods).
- Stream attribute and route algebra are just **higher abstractions** of live variables.
- The update of live variables satisfies the **glitch-free consistency**.

Glitch-freedom⁴: Intermediate consequences of a data change must not be observed.



⁴Alessandro Margara and Guido Salvaneschi. "We Have a DREAM: Distributed Reactive Programming with Consistency Guarantees". In: DEBS '14. New York, NY, USA: ACM, 2014.

TIE EVERYTHING TOGETHER

Example Program: Block traffic for infected hosts and construct routes using concatenation

Events:

```
1  iff (pkt.is_endhost_infected) {  
2    drop(pkt)  
3  } else {  
4    bind(pkt, r_1 + r_2)  
5  }
```




TIE EVERYTHING TOGETHER

Example Program: Block traffic for infected hosts and construct routes using concatenation

Events:

- When any network component on r_1/r_2 changes, the new route ($r_1 + r_2$) is automatically recomputed

```
1  iff (pkt.is_endhost_infected) {  
2      drop(pkt)  
3  } else {  
4      bind(pkt, r_1 + r_2)  
5  }
```




TIE EVERYTHING TOGETHER

Example Program: Block traffic for infected hosts and construct routes using concatenation

Events:

- When any network component on r_1/r_2 changes, the new route ($r_1 + r_2$) is automatically recomputed
- **When the host status changes to infected, all packets are dropped**

```
1  iff (pkt.is_endhost_infected) {  
2      drop(pkt)  
3  } else {  
4      bind(pkt, r_1 + r_2)  
5  }
```




TIE EVERYTHING TOGETHER

Example Program: Block traffic for infected hosts and construct routes using concatenation

Events:

- When any network component on r_1/r_2 changes, the new route ($r_1 + r_2$) is automatically recomputed
- When the host status changes to infected, all packets are dropped
- When the host status is cleared (e.g., through an admin interface or timeout), a route $r_1 + r_2$ is automatically recomputed

```
1  iff (pkt.is_endhost_infected) {  
2      drop(pkt)  
3  } else {  
4      bind(pkt, r_1 + r_2)  
5  }
```



IMPLEMENTATION HIGHLIGHTS

Efficient update: *how does Trident achieve fast updates?*

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```


- Assume a change to G that invalidates both p and b .
- Trident recomputes p and b .
- Trident returns r as soon as p is ready and not `Unknown`.

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```




At time 0: no routes is ready

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```




At time 1: the primary route $p0$ is ready. With efficient update, Trident selects $p0$ as r .

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```




At time 2: the backup route $b0$ is also ready. The standard update will select $p0$ as r .

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```




At time 3: a data change e happens and invalidates the primary route.

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```




At time 3: now p is not ready, Trident selects $b0$ as r .

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```




At time 4: p has a new value $p1$. With efficient update, Trident selects $p1$ as r .

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

Example:

```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```



With this simple code, Trident achieves lifecycle management for backup routes.

EFFICIENT UPDATE

Objective: Leverage semantics of operators to achieve fast updates.

General rule: For all route algebra operators, if the partial result has no unknown subsets, the output guarantees glitch-free consistency and we can apply efficient update.

<i>Expr</i>	<i>Known Subset</i>	<i>Unknown Subset</i>
$\Delta_1 \cup \Delta_2$	$\mathcal{K}_1 \cup \mathcal{K}_2$	$\mathcal{U}_1 \cup \mathcal{U}_2$
$\Delta_1 \cap \Delta_2$	$\mathcal{K}_1 \cap \mathcal{K}_2$	$(\mathcal{K}_1 \cap \mathcal{U}_2) \cup (\mathcal{U}_1 \cap \mathcal{K}_2) \cup (\mathcal{U}_1 \cap \mathcal{U}_2)$
$\Delta_1 \setminus \Delta_2$	$T_{\mathcal{U}_2=\emptyset}(\mathcal{K}_1 - \mathcal{K}_2)$	$(T_{\neg(\mathcal{U}_2=\emptyset)}(\mathcal{K}_1) \cup \mathcal{U}_1) - (\mathcal{K}_2 \cup \mathcal{U}_2)$
$\Delta_1 + \Delta_2$	$\mathcal{K}_1 + \mathcal{K}_2$	$(\mathcal{K}_1 + \mathcal{U}_2) \cup (\mathcal{U}_1 + \mathcal{K}_2) \cup (\mathcal{U}_1 + \mathcal{U}_2)$
$\asymp \Delta$	$\asymp \mathcal{K}$	$\asymp \mathcal{U}$
$\sigma_f(\Delta)$	$\sigma_f(\mathcal{K})$	$\sigma_f(\mathcal{U})$
$\diamond_d(\Delta)$	$T_{\mathcal{U}=\emptyset}(\diamond_d(\mathcal{K}))$	$\diamond_d(T_{\neg(\mathcal{U}=\emptyset)}(\diamond_d(\mathcal{K})) \cup \diamond_d(\mathcal{U}))$
$\Delta_1 \triangleright \Delta_2$	$\mathcal{K}_1 \cup T_{\mathcal{U}_1=\emptyset}(\mathcal{K}_2 - \mathcal{K}_1)$	$\mathcal{U}_1 \cup ((T_{\neg(\mathcal{U}_1=\emptyset)}(\mathcal{K}_2) \cup \mathcal{U}_2) \setminus (\mathcal{K}_1 \cup \mathcal{U}_1))$
$*\Delta$	$*\mathcal{K}$	$T_{\mathcal{K}=\emptyset}(*\mathcal{U})$

$T_\varepsilon(S)$ - the value is $S \cup \{\varepsilon\}$ if $\varepsilon = \text{true}$, and $\{\varepsilon\}$ otherwise.

Table 1: Known/Unknown Subsets of Route Algebra.

EVALUATION

- How much effort does one need to integrate a network function into Trident?
- How useful is efficient update?

EVALUATION SETTINGS

- CPU: Intel Xeon CPU E5-2650 2.30GHz
- Memory: 64G
- OS: Fedora 26
- Network: Mininet 2.3.0d1

MODIFICATION ON NETWORK FUNCTIONS

Two case studies:

- Bro: a deep packet inspection framework
- FreeRadius: an open-source Radius server

<i>Name</i>	<i>Attribute</i>	<i>Language</i>	<i>LoC (f)</i>	<i>LoC (a)</i>	<i>LoC (c)</i>
DPI	HTTP URL	Bro	40	2	2
FreeRadius	Auth status	DSL	0	12	0

LoC - Additional lines of code, **f** - LoC to implement the library in the given framework/language, **a** - LoC in a given NF, **c** - LoC for configuration.

EFFICIENT UPDATE MICRO BENCHMARK

We demonstrate the effect of efficient update by measuring the recovery time of


```
val p = ShortestPath(G, s, t)
val ps = snapshot(p)
val b = ffr(G, ps)
val r = any(ps >> b)
```

for a single (src, dst) pair in 4 different topologies.


We compare the results of

- The initial computation (**Init**) v.s. fast rerouting stage (**FR**)
- The standard update method (**SCP**) v.s. the efficient update (**ECP**)


EFFICIENT UPDATE MICRO BENCHMARK




Sprint (11 nodes)



Noel (19 nodes)




Agis (25 nodes)




Geant (40 nodes)


EFFICIENT UPDATE MICRO BENCHMARK




Sprint (11 nodes)



Noel (19 nodes)



Agis (25 nodes)



Geant (40 nodes)

Trident demonstrates an improvement of up to 39% in initial computation and an improvement of 1 to 2 magnitudes in the fast rerouting stage.

SUMMARY

Trident is a unified SDN programming framework which

- uses **stream attribute** to naturally integrate network function state into logically centralized SDN programming;
- uses **route algebra**, a simple yet powerful abstraction to systematically construct consistent, correlated routes;
- uses **live variable** to achieve unified automatic data dependency management and glitch-free updates – achieving a more general **intent** networking framework.

Future directions:

- Extend from “write-only” network functions to generic network functions
- Verification with Trident

THANKS FOR YOUR ATTENTION!

Q & A

KAI GAO
GODRICKK@GMAIL.COM

TAISHI NOJIMA
TAISHI.NOJIMA@AYA.YALE.EDU

Y. RICHARD YANG
YRY@CS.YALE.EDU

DESIGN SPACE

Basic Programming Model

SDN Programming Languages

Define behaviors as if processing **each packet** based on its attributes

- Simple programming paradigm
- Cannot handle layer-7 information naturally

Event-Driven SDN+NF Systems

Define behaviors by specifying **state** and **transitions**

- Fit well with how NF process packets
- Require *manual* efforts to identify transitions and can be complex when many NF states are involved

Can we inherit the simple paradigm of SDN programming but integrate network function states naturally?

DESIGN SPACE

Route Specification

Fully Customizable Routing

The policy can execute a routing function over the graph to find the path

- Very flexible
- Complex
- May not be efficient

One-Big-Switch Abstraction

The policy only specifies the output logical port

- Simple
- Optimized by the system
- Not flexible
- Cannot express path requirements

Can we have a simple abstraction for route specification but still retain enough flexibility?

DESIGN SPACE

Handling Dynamicity

Pre-defined Dependencies

The system predefines some dependencies with domain knowledge (such as failure model of network components)

- Automatic
- Simplify the job for programmers
- Hard-coded


Manual Dependency Management

Programmers manually identify and handle all dependencies (for example, using the Observer pattern)

- Optimized
- Accurate
- Error-prone
- Difficult to manage and maintain

Can we achieve automatic dependency management with guaranteed correctness (i.e., data plane configuration is consistent with the control plane state)?

MOTIVATING EXAMPLE



- HTTP connections with a trusted URL **can** skip the DPI nodes
- HTTP connections with a trusted URL to a large-scale data transfer service **should** use high-bandwidth links whenever possible

Why cannot existing SDN programming languages work?

The ambiguity of “failed” predicates with network function states

- For `sip == 10.0.0.2`, the result is either true or false for all packets
- For `is_trusted(http_url)`, the result can be true, false and **unknown**

Why cannot existing SDN programming languages work?

The correlation of routes cannot be explicitly specified

- Some network functions require traffic of the same stream to be processed on the same node and must migrate to the same new instance simultaneously (e.g., TCP & HTTP state machine)
- Some routes depend on others (e.g., link protection)
- Some routes depend on certain conditions (e.g., link capacity requirement)

Why cannot existing SDN programming languages work?

The dynamicity in both packet selection and route computation is closely related

- A single NF state may affect multiple streams (for example, streams with the same heavy hitter status may have different HTTP URL values)
- Routes for multiple streams are computed together by a single routing algorithm