

NFD: Using Behavior Models to Develop Cross-Platform Network Functions

Hongyi Huang, Wenfei Wu, Yongchao He, Bangwen Deng, Ying Zhang,
Yongqiang Xiong, Guo Chen, Yong Cui, Peng Cheng

Tsinghua University, Facebook, Microsoft Research, Hunan University

Content



01

Background

02

Goal&Challenge

03

Design

04

Evaluation

05

Conclusion



Content



01

Background

02

Goal&Challenge

03

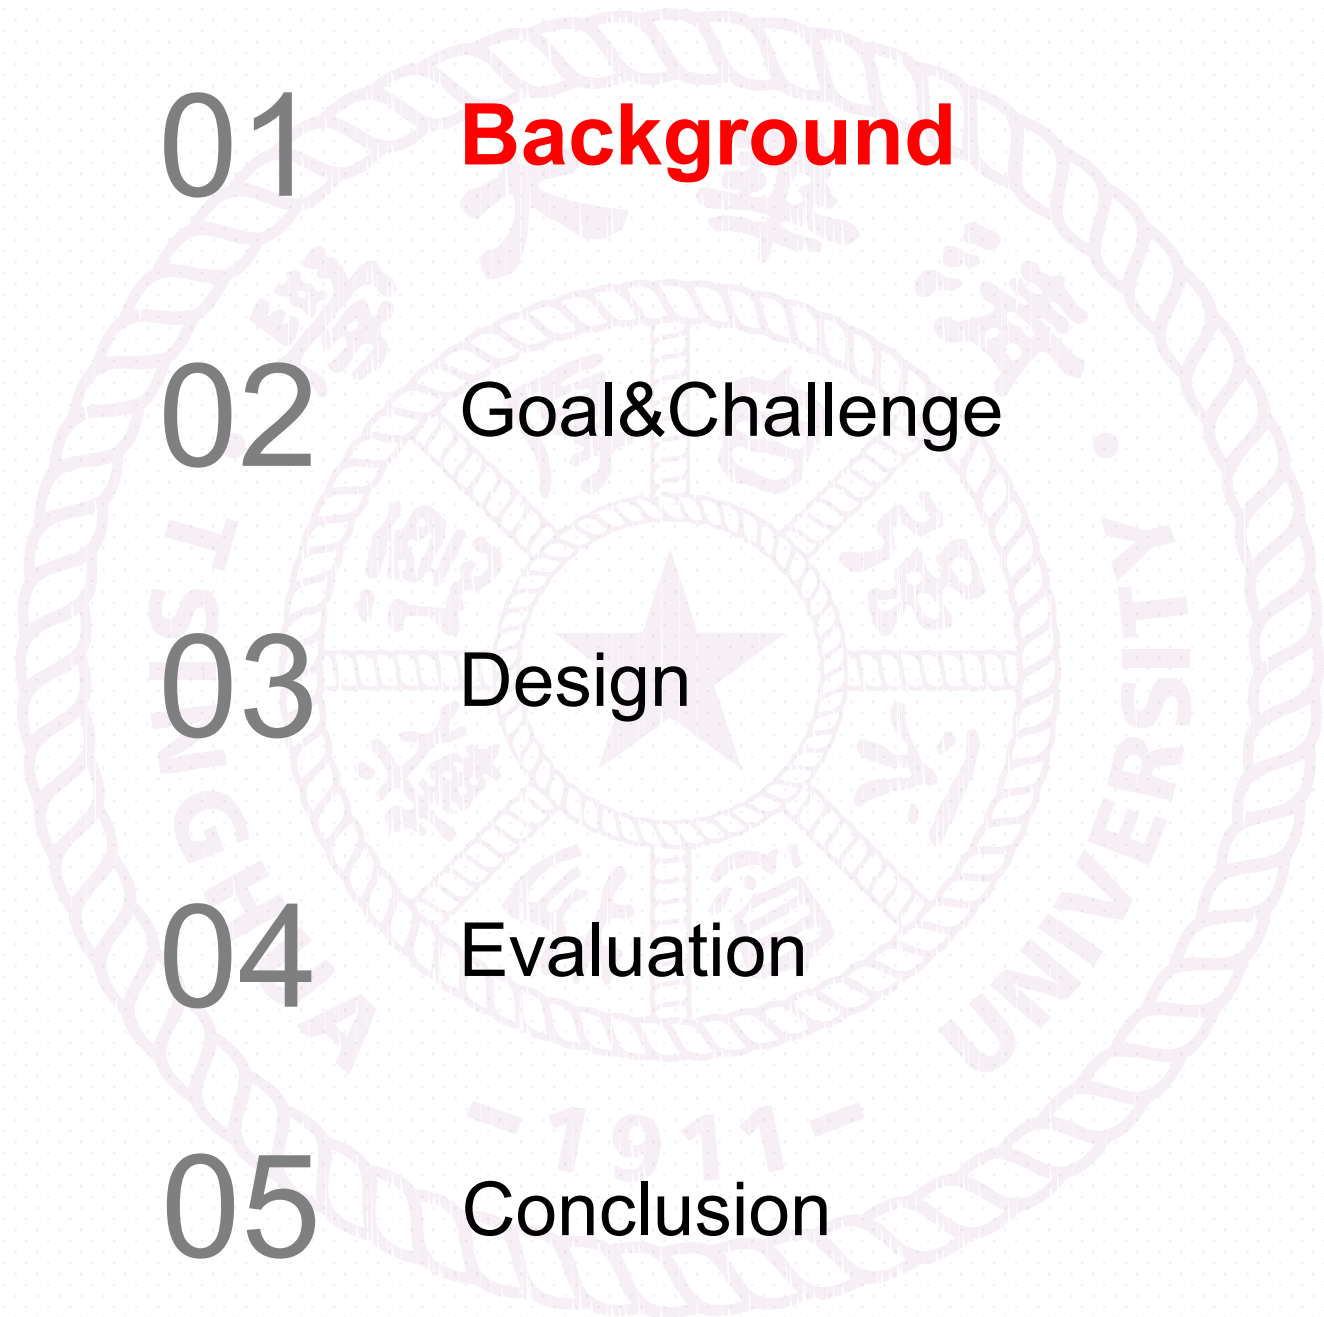
Design

04

Evaluation

05

Conclusion



Network Function (NF) / Middlebox

Various network functions are widely deployed in network & between hosts in addition to switches & routers

- hard-coded
- wired



Wan optimizer



Proxy



IDS

.....



CDN



Firewall



NAT

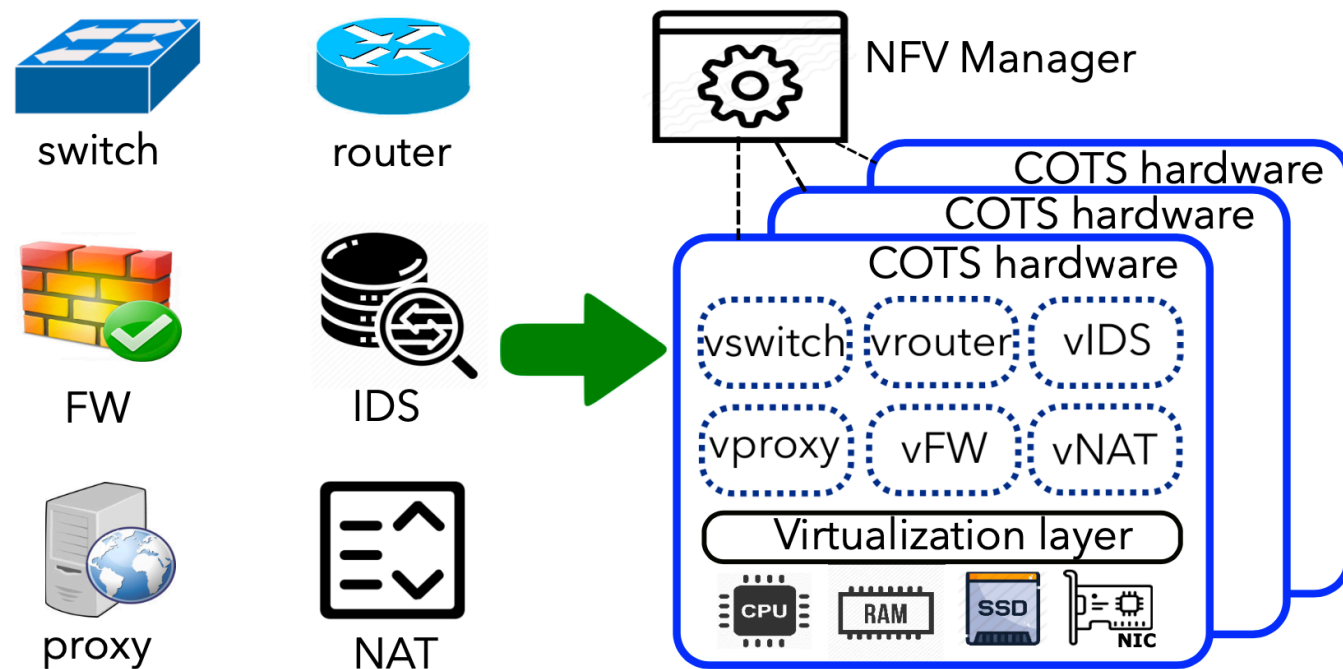


BRAS

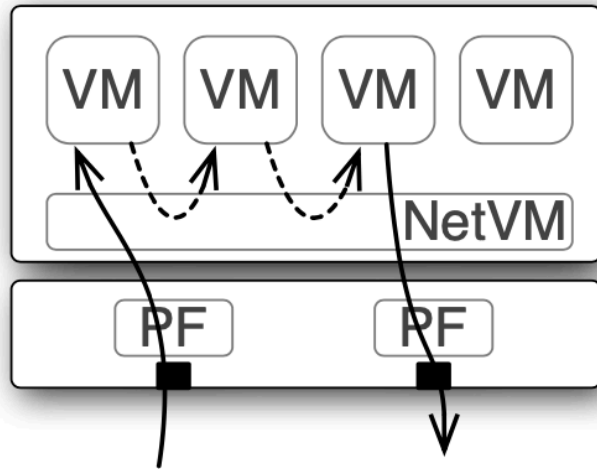
NFV

NFV replaces specialized middleboxes with software Virtual Network Functions (VNFs) consolidated on Commodity Off-The-Shelf (COTS) hardware.

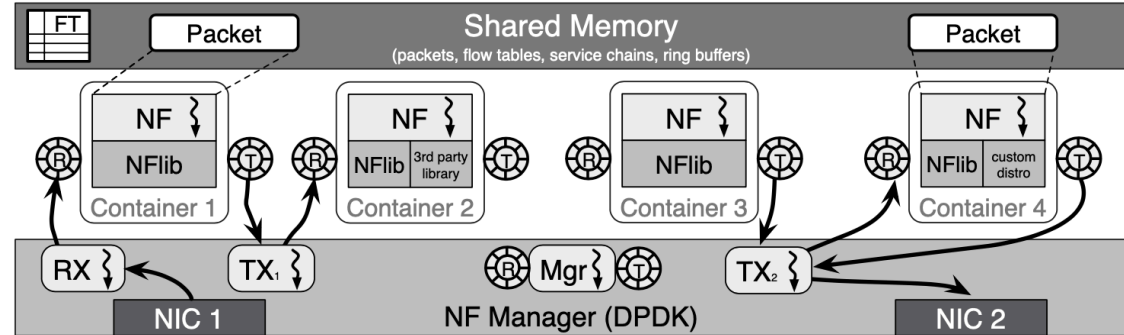
- Flexible deployment
- Quick evolution



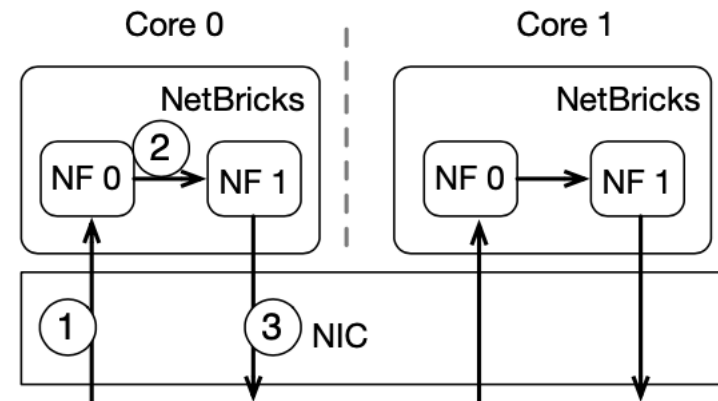
Deployment Way



Virtual Machine: NetVM, ClickOS

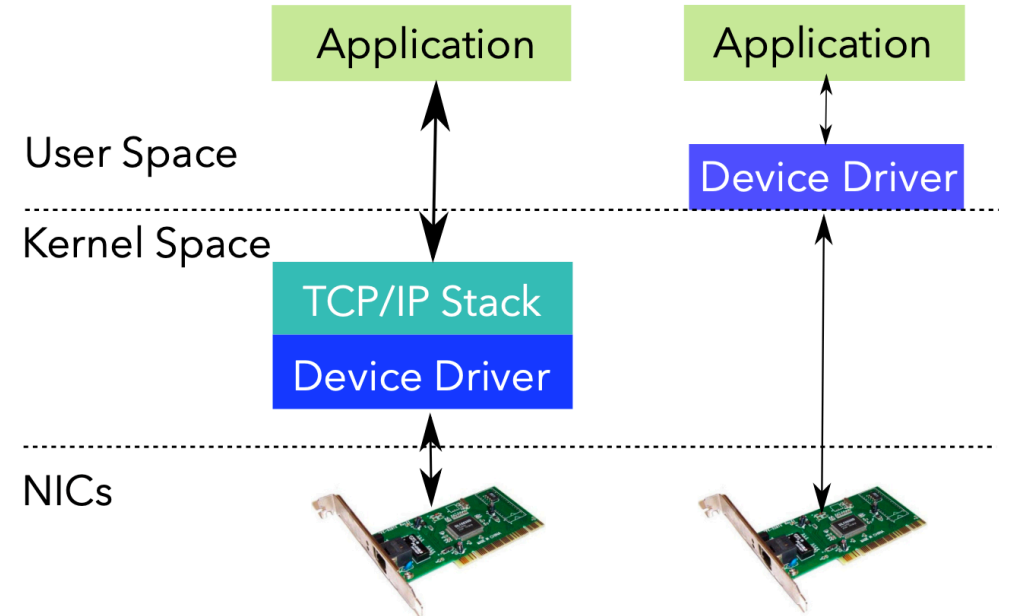
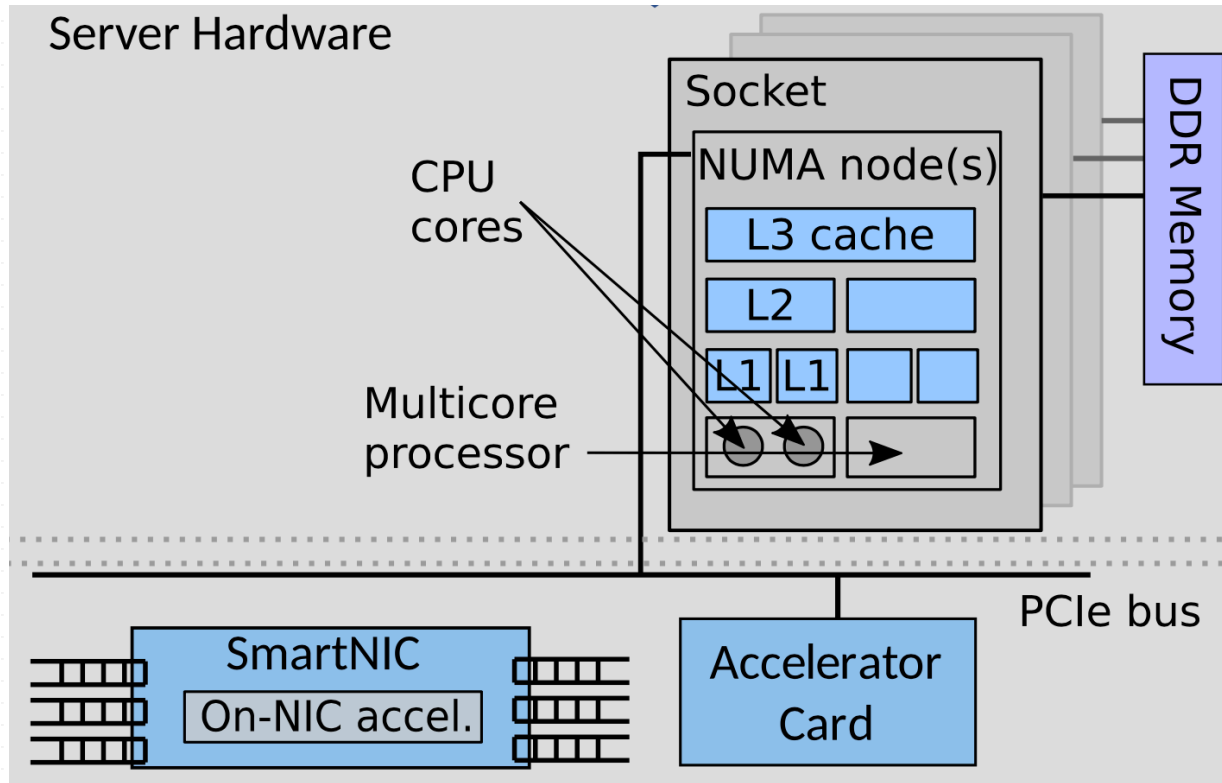


Container: OpenNetVM



Process/Function: Click, NetBricks

Acceleration Technique



Hardware-assisted: Floem, ClickNP, ResQ, SGX(Sgx-box), DPDK(many), GPU(apunet)

Software: OpenNF, S6

Integration: Cross-Product

NF Developer

IDS(Intrusion Detection System)
NAT(Network Address Translator)
Load Balancer
Rate Limiter
Cache
Monitor
...
...

m

Platform Provider /Developer

NetVM
OpenNetVM
Click
NetBricks
OpenNF
...
...

n

Network Operator



Explosion

$m \cdot n$

X

=

IF $m+n$ Possible?

m NFs + n Platforms(Compiler)

→ $m * n$ (integrations)

m NF models + n Platform(Abstraction + Compiler)

→ $m * n$ (models) + n compilers
auto integration

m NF models + 1 Framework(Abstraction + Compiler) + n Platform Plugins

→ m (models) + n (plugins) + 1 Framework
auto integration

Content



01

Background

02

Goal&Challenge

03

Design

04

Evaluation

05

Conclusion



Example 1

Kernel bypassing I/O

Change packet I/O with DPDK-enabled NICs

original

```
//=== Snort.c with libpcap ===  
int main(int argc, char* argv[]) {  
    ... // initialization  
    pcap_loop(phandle, -1, pcap_handle, NULL);  
}
```

modified

```
//=== Snort.c with DPDK ===  
int main(int argc, char* argv[]) {  
    ... // initialization  
    for (;;) {  
        struct rte_mbuf *bufs[SIZE];  
        rte_eth_rx_burst(port, 0, bufs, SIZE);  
    }  
}
```

Example2

State migration and management

Modifying PRADS and Snort to integrate with **OpenNF** takes more than 100 man-hours
[OpenNF, StateAlyzr]

```
//=== Snort.c without OpenNF ===  
int main(int argc, char* argv[]){  
    ... // initialization  
    pcap_loop(phandle, -1, pcap_handle, NULL);  
}
```

```
//=== Snort.c with OpenNF ===  
int main(int argc, char* argv[]){  
    ... // initialization  
    locals.put_allflows = &put_allflows;  
    sdmbn_init(&locals); // start agent  
    pcap_loop(phandle, -1, pcap_handle, NULL);  
}
```

Example3

secure VNFs from memory reading attacks

2.5k extra lines of code in the modification when porting an IDS to **Intel SGX** [SGX-box]

state

```
//=== PRADS.c ===  
void check_vlan (packetinfo *pi) {  
    config.pr_s.vlan_recv++; // a state  
    ...  
}
```

state

```
void prepare_ip4 (packetinfo *pi){  
    config.pr_s.ip4_recv++; // a state  
}
```

```
//=== SGX Config ===  
enclave {  
    ...  
    trusted{ public void check_vlan (...);  
             public void prepare_ip4 (...);  
};};
```

Intuitions to Design New Framework

- a). Most integration targets a **specific piece of logic**(e.g., IO in DPDK, states in OpenNF)
- b). The framework should provide **interfaces** for *logic identification* in NFs.
- c). Integration operations are tedious but **regular**.

Goal: Build a cross-platform development framework for NFs

Challenge

1. Expressiveness.

Expressive to describe the packet processing logic in various NFs

2. EasyDev.

How to save the development workload for platform provider

3. Performance.

The outcome NFs should be logically correct and have comparable performance compared with existing legacy NFs.

Content



01 Background

02 Goal&Challenge

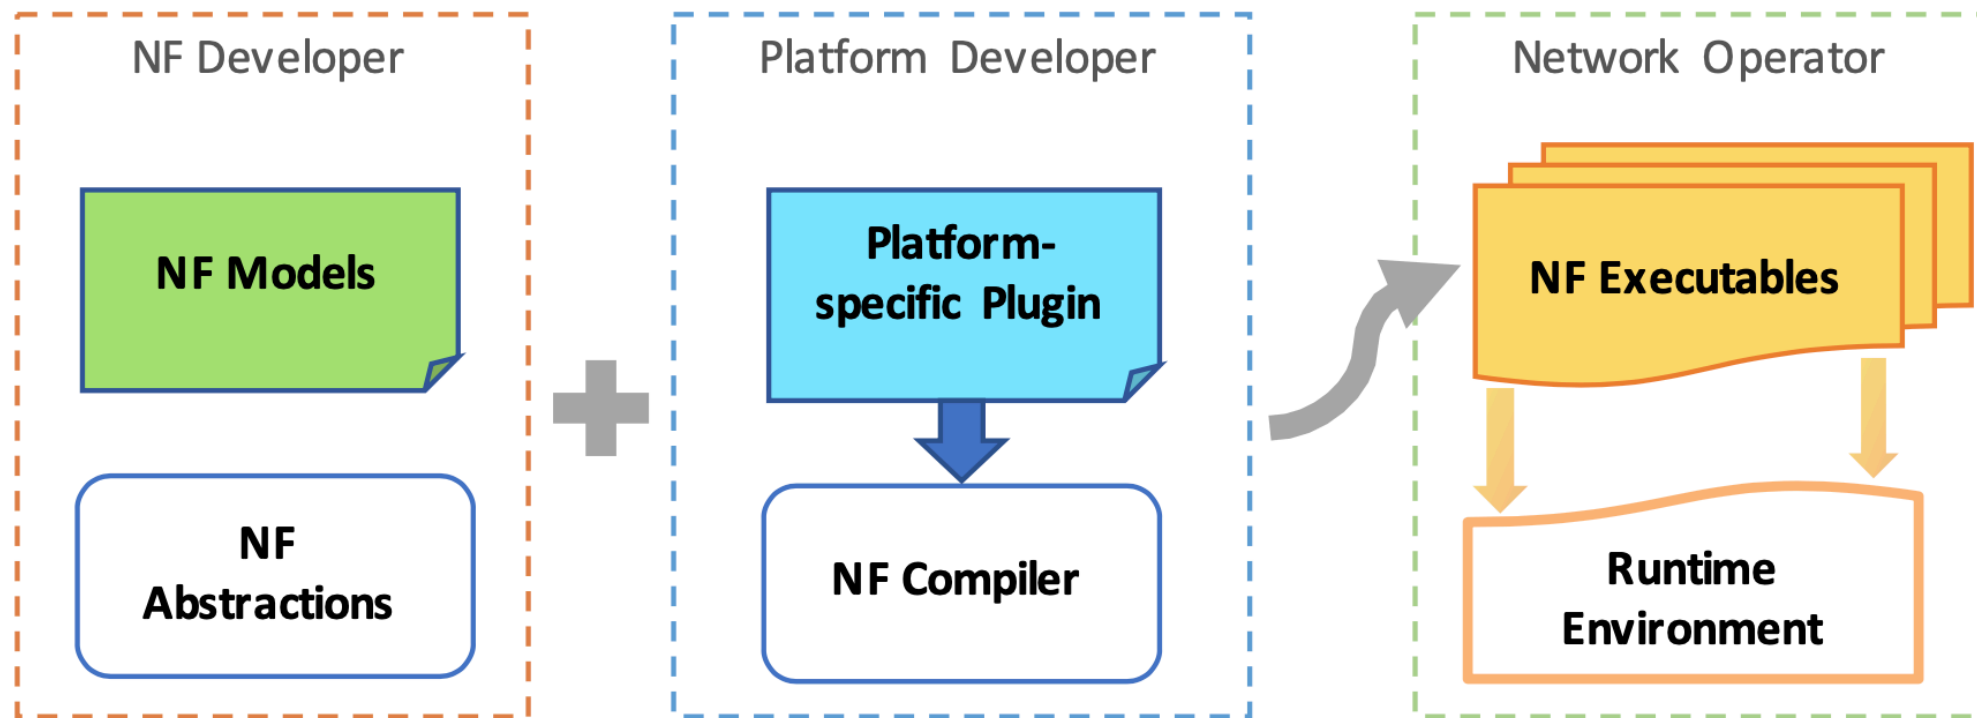
03 **Design**

04 Evaluation

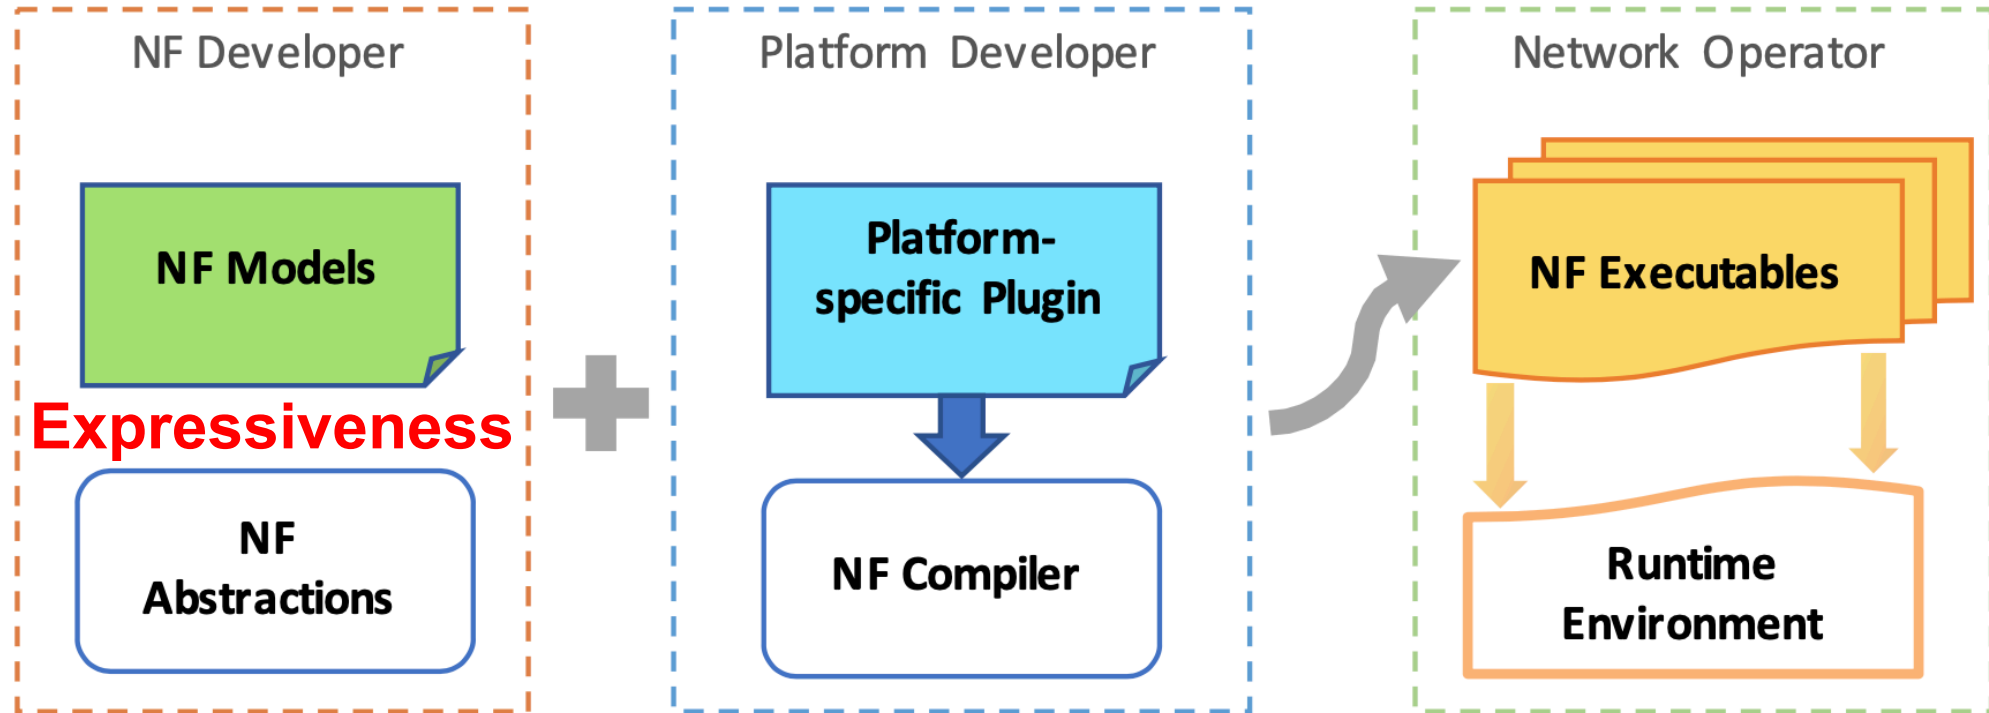
05 Conclusion



Overview



Overview



Expressive Language

Existing programming frameworks:

NetCore, SNAP, SDN, Click, ...

Basic types and expression

const	c	$::=$	$(0 1)^+$
header field	h	$::=$	$sip dip sport dport proto ...$
state	s		
variable	var		
expression	e	$::=$	$c h s var e Expr_Op(e_1, e_2, ...)$

Predicates

flow predicate	x_f, y_f	$::=$	$\epsilon * h = c \neg x_f x_f \wedge y_f x_f \vee y_f$
state predicate	x_s, y_s	$::=$	$* Rel_Op(s, e) \neg x_s x_s \wedge y_s x_s \vee y_s$
general predicate	x, y	$::=$	$Rel_Op(e_1, e_2, ...) \neg x x \wedge y x \vee y$

Policies and Statements

flow policy	p_f, q_f	$::=$	$h := e p_f ; q_f$
state policy	p_s, q_s	$::=$	$s := e p_s ; q_s$
general policy	p, q	$::=$	$q := e p ; q$

Model

model	$model$	$::=$	$stmts$
statements	$stmts$	$::=$	$stmt stmt ; stmts$
statement	$stmt$	$::=$	$p if loop$
if statement	if	$::=$	$if(x) \{ stmts \} else \{ stmts \}$
loop statement	$loop$	$::=$	$while(x) then \{ stmts \}$

Semantics

symbols	meaning
$f[h]$	h is a header field (Figure 4 does not list all fields), and $f[h]$ is the field h in packet f .
$f[TAG]$	We append tags to each packet for flexible processing[34], which can be viewed fields of a packet.
$f[output]$	Record the output ports of a packet. $f[output] := \{p_1, p_2\}$ means sending packet f to port p_1 and p_2 . $f[output] := \epsilon$ means dropping the packet.
r	Abbreviation for A rule: $h_1 = v_1 \wedge h_2 = v_2 \wedge \dots$
$f \in r$	Abbr. for a flow-rule match: $f[h_1] = v_1 \wedge f[h_2] = v_2 \wedge \dots$
R	Abbreviation for a rule set: $\{r_1, r_2, \dots\}$
$f \in R$	Abbreviation for a flow-ruleset match (f match one of rules in R): $f \in r_1 \vee f \in r_2 \vee \dots$

Programming Abstractions

Packet processing abstraction: parse, deparse, transform

Bytestream processing abstraction: TCP flow

User-defined abstraction: custom abstractions

State abstraction: manage state in granularity

```
1  string type="int";  Value value=0;
2  int granularity=sip&dip&sport&dport&proto;
3  map<unsigned, Value> instances;
4  State_Counter& operator++(){
5      key=hash(pkt&MaskOf(granularity));
6      if(instances.find(key)==instances.end())
7          instances.put(key, value);
8      instances[key]++;
9  }  };
```

Time-driven logic abstraction: timer

Uniform Structure

Stateful Match Action Table

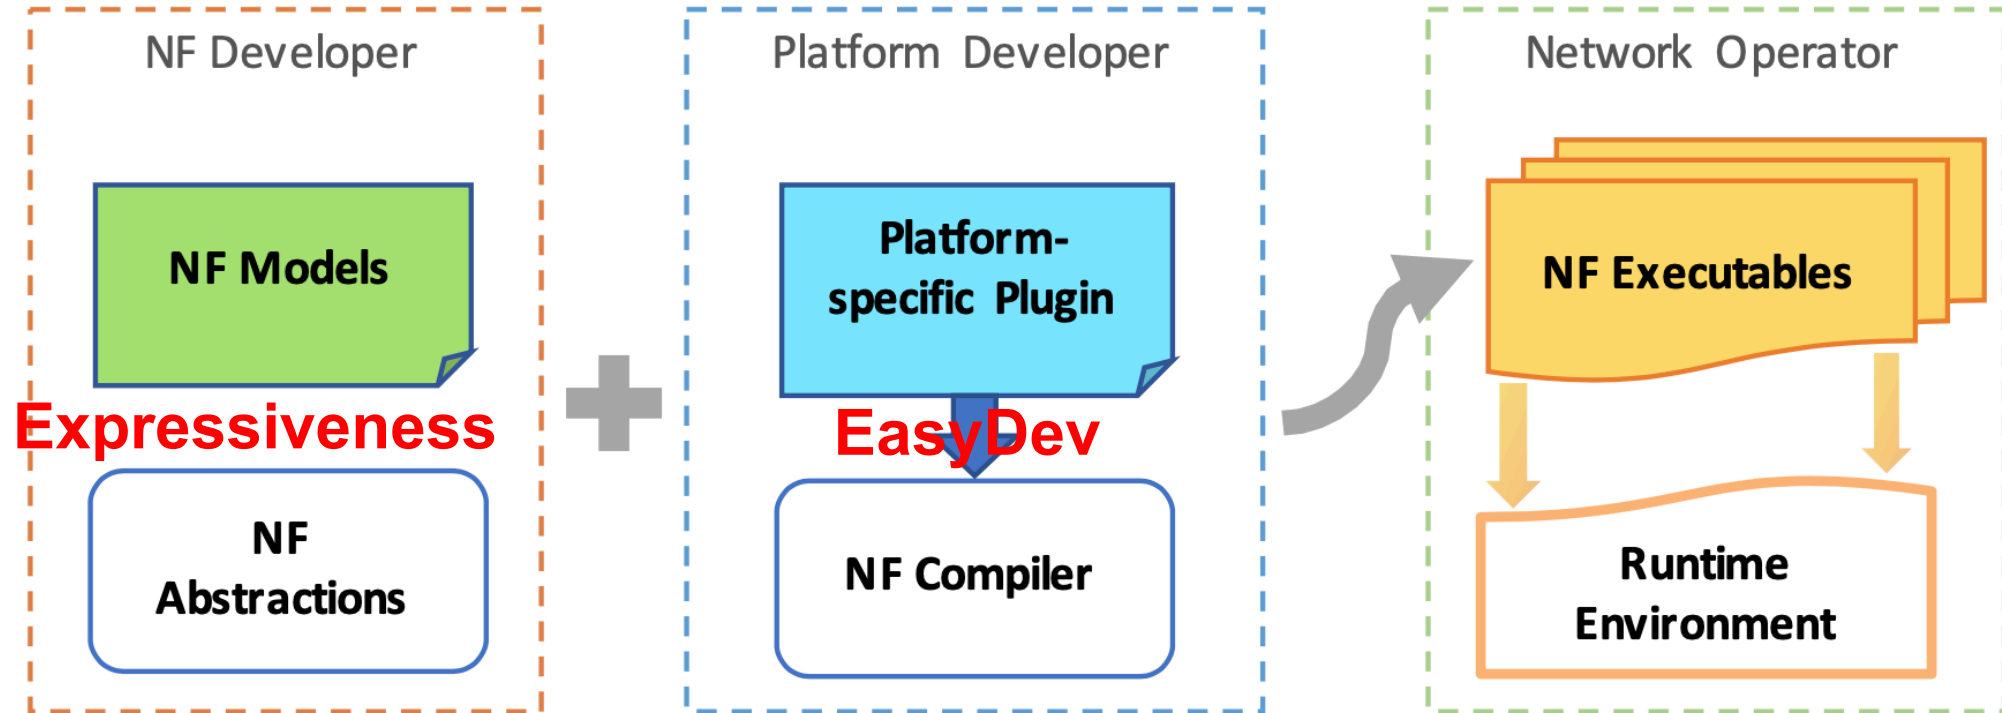
$\text{entry } \text{entry} ::= \text{if } (x_f \wedge x_s) \text{ then } (p_f; p_s) \text{ else } \perp$
 $\text{SMAT } \text{smat} ::= \text{entry} | \text{entry}; \text{model}$

Rationality:

1. Existing practices in Microsoft Azure [VFP]
2. Stateless -> compatible with switch policy [SDN, P4]

	Match		Action	
	Flow	State	Flow	State
Stateful Firewall	Configuration: $OK = \{r1, r2, \dots\}$			
	$f \in OK$	-	$f[\text{output}] := \text{IFACE}$	$\text{seen} := \text{seen} \cup \{f\}$
	f	$f \in \text{seen}$	$f[\text{output}] := \text{IFACE}$	-
	$f \notin OK$	$f \notin \text{seen}$	$f[\text{output}] := \epsilon$	-

Overview

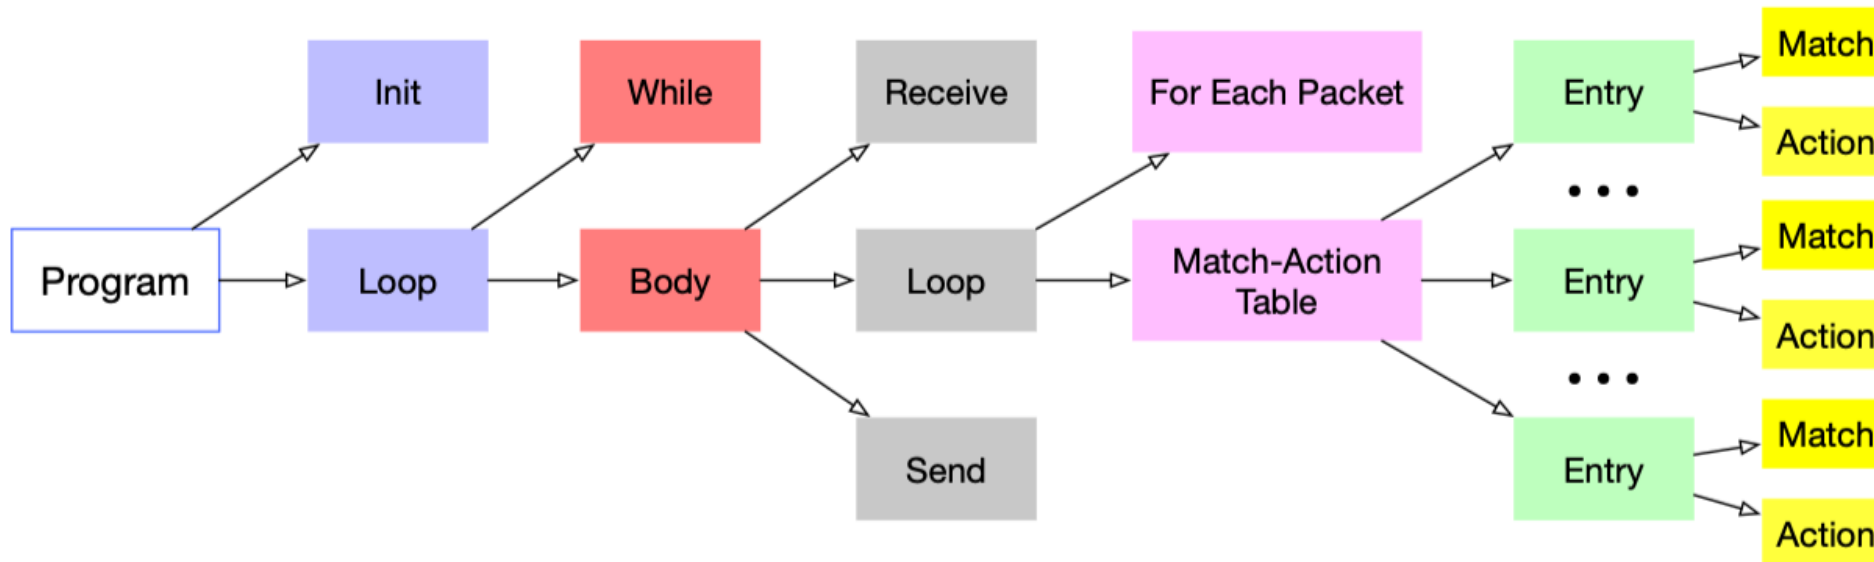


Syntax Tree

Intermediate Representation (IR)

leaf nodes

non-leaf nodes : derived to leaf nodes



EasyDev.

Interface1: Override (DPDK, GPU)

-- replace a piece of logic by a platform enhanced implementation
(replace pcap_loop)

```
//=== Snort.c with libpcap ===  
int main(int argc, char* argv[]){  
    ... // initialization  
    pcap_loop(phandle, -1, pcap_handle, NULL);  
}
```

```
//=== Snort.c with DPDK ===  
int main(int argc, char* argv[]){  
    // initialization  
    for (;;) {  
        struct rte_mbuf *bufs[SIZE];  
        rte_eth_rx_burst(port, 0, bufs, SIZE);  
        ...  
    }  
}
```

EasyDev.

Interface2: Modification (OpenNF)

-- insert/delete/modify a node on the syntax tree using IR callback function
(add initialization)

```
1   new OpenNFVisitor.visit(syntax_tree);
2   }
3   public class OpenNFVisitor implements NFDCompiler{
4   @Override public T visitInit(...){
5       AddAgentCode(...)
6       InsertCode("List<State>_allStates")
7       super.visitInit(...) // orig. compilation }
8   @Override public T visitStateDeclaration(...){
9       super.visitStateDeclaration(...)
10      stateName = ... // get the state name
11      InsertCode(String.format("allStates.add(%s)", stateName))
12  } }
```

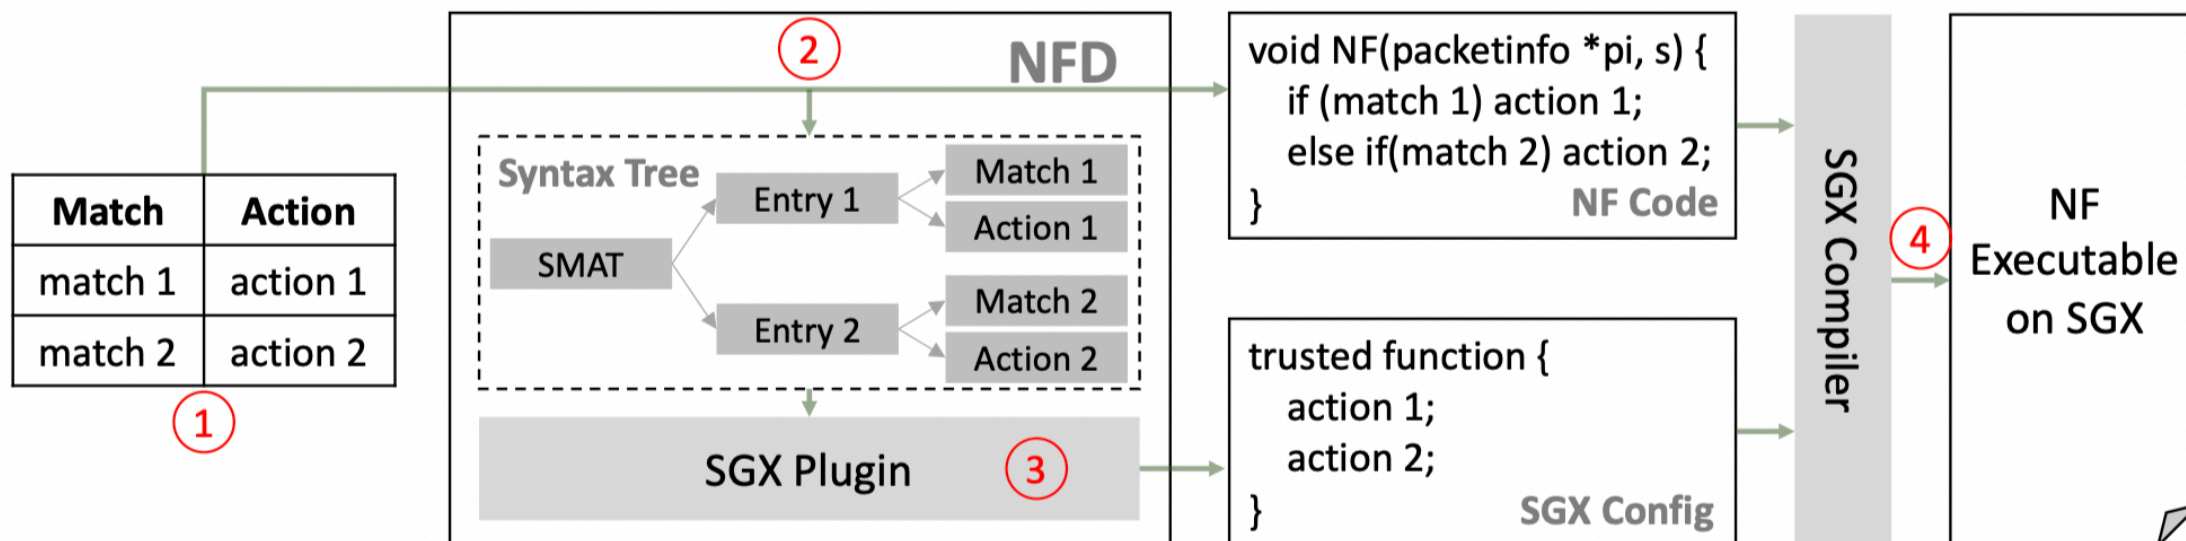
EasyDev.

Interface3: Retrieval (SGX)

- collect extra information (sensitive state and function)
- use the information for platform integration (SGX config)

```
13     new SGXVisitor.visit(syntax_tree);
14 }
15 public class SGXVisitor implements NFDCompiler{
16     List<String> sensitiveFunc;
17     List<String> sensitiveData;
18     @Override public T visitStateDeclaration (...) {
19         stateName =
20         sensitiveData.add(stateName);}
21     @Override public T visitStateMatch (...) {
22         FuncName = ...
23         sensitiveFunc.add(FuncName);}
24     @Override public T visitStateAction (...) {
25         FuncName = ...
26         sensitiveFunc.add(FuncName);}
27 }
```

Example Workflow



Content



01 Background

02 Goal&Challenge

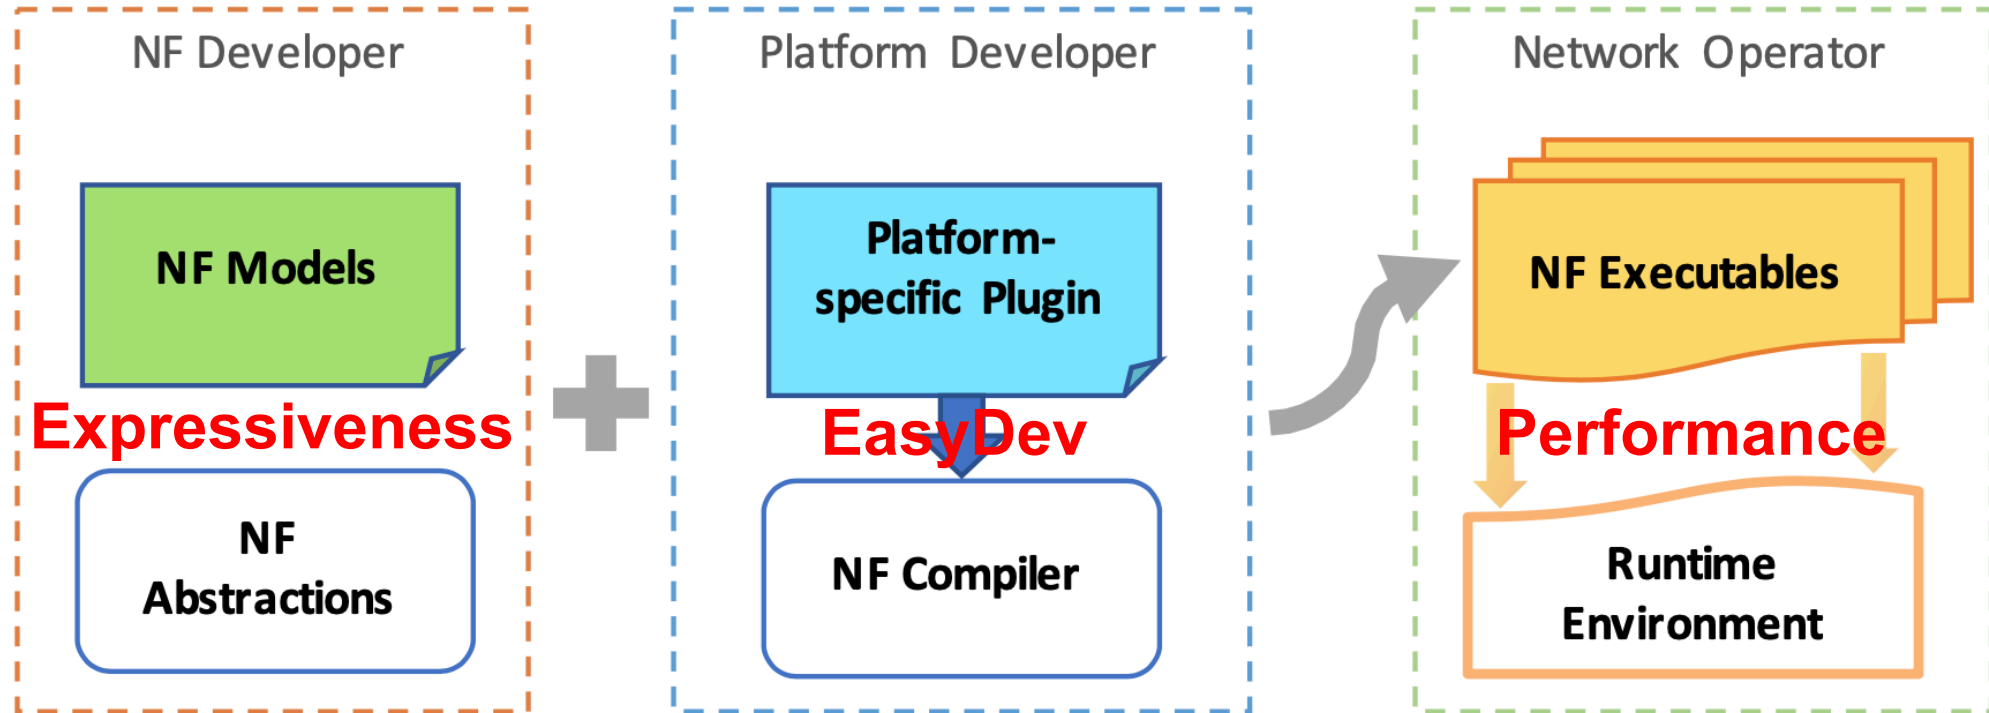
03 Design

04 **Evaluation**

05 Conclusion



Overview



Prototype

Source: DSL
 Compiler: ANTLR4
 Target: C++

Component of NFD	Lines of Code
NFD model grammar	234 (g4)
compiler frontend (automatically derived by Antlr)	4.3k (Java)
compiler backend (generate C++ NF programs)	1137 (Java)
C++ template (program structure, operators) for NFs	752 (C++)
extension for OpenNF	489 (C++)
extension for GPU	668 (C++)
extension for DPDK	167 (C++)
extension for SGX	273 (C++)

14 NFs + 6 Platforms

Comparing Workload (LOC)

NFD: models + framework + plugins = ~ 4k

manual: ~700k

Testbed

Server:

- Intel i9 CPU (10-core, 20-thread)
- 128GB memory
- 10Gbps NIC
- three NVIDIA GTX1080 Ti graphics cards
- 1TB SSD

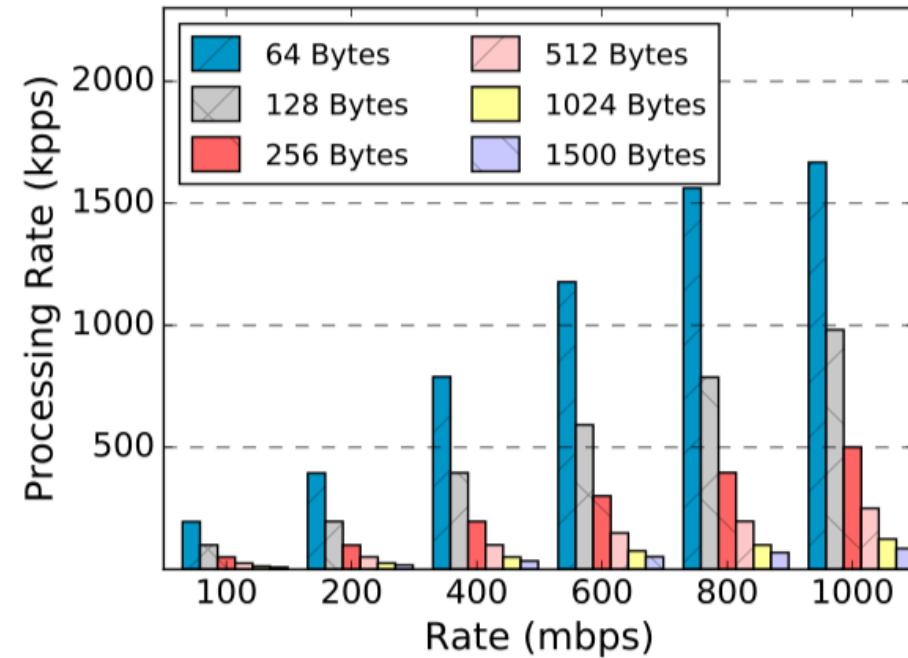
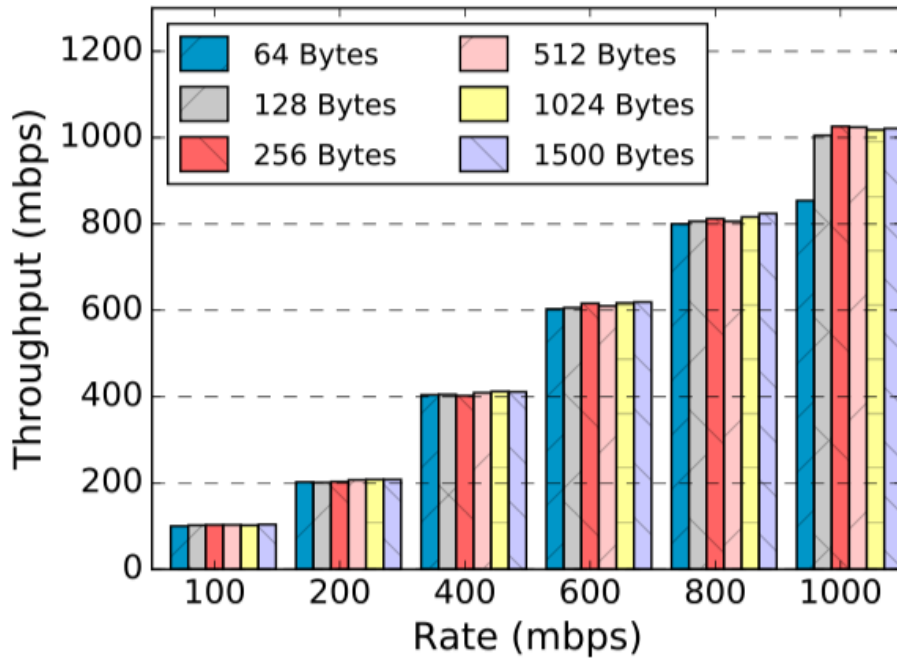
Trace: [IMC10]

Opensource NFs: Snort, PRADS, Balance, HAProxy, Click NAT

Correctness

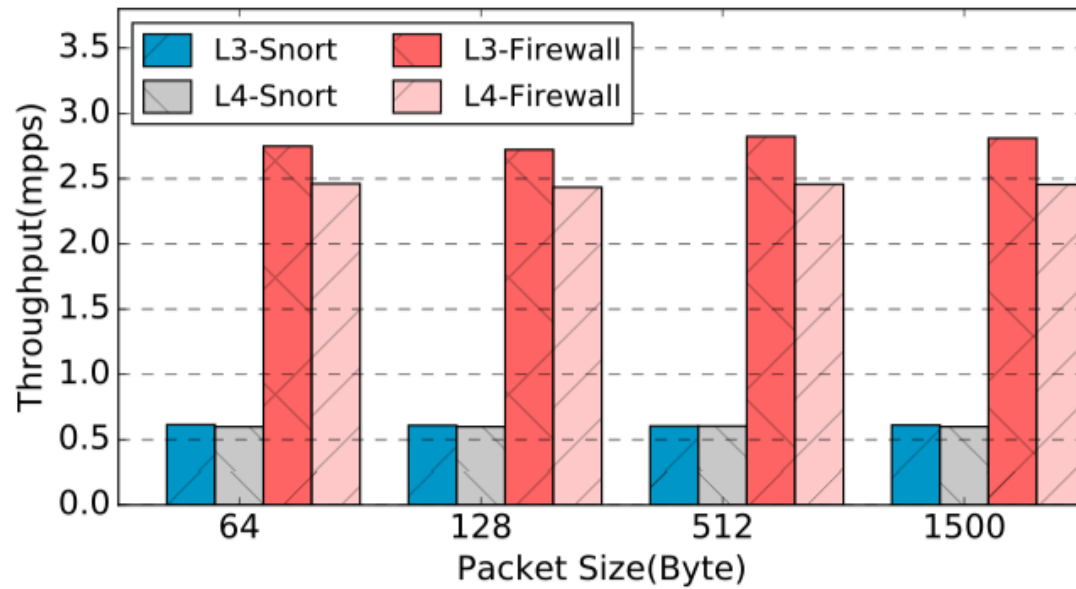
Rate Limiter: tuning rate

-- control the sending rate accurately as configuration



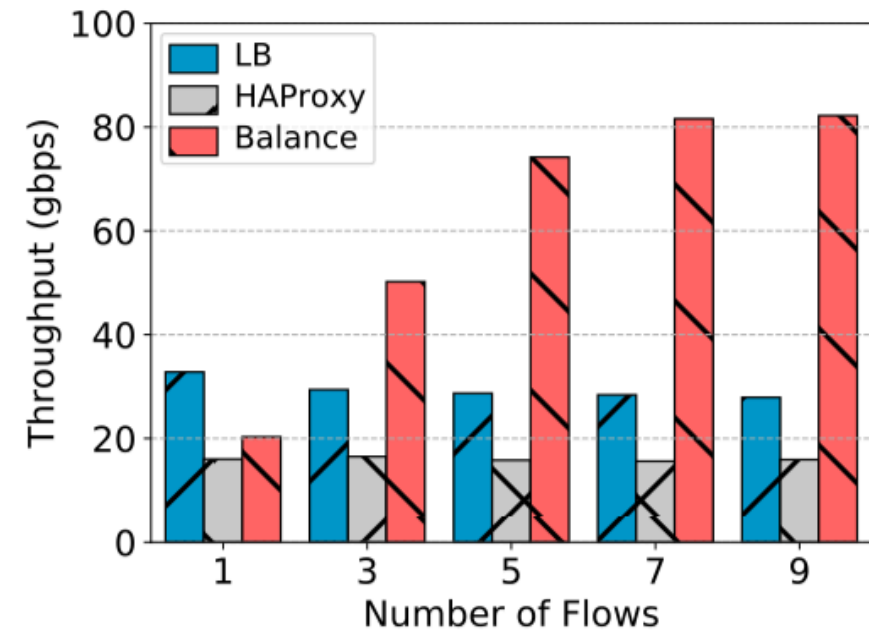
Performance

Firewall



Optimization: reduce redundant logic

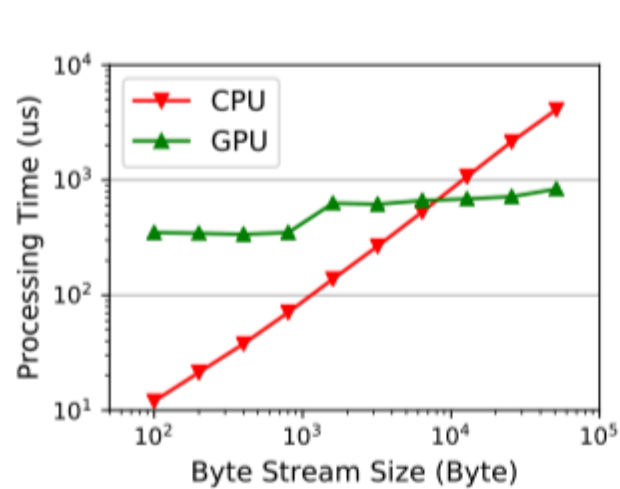
Load Balancer



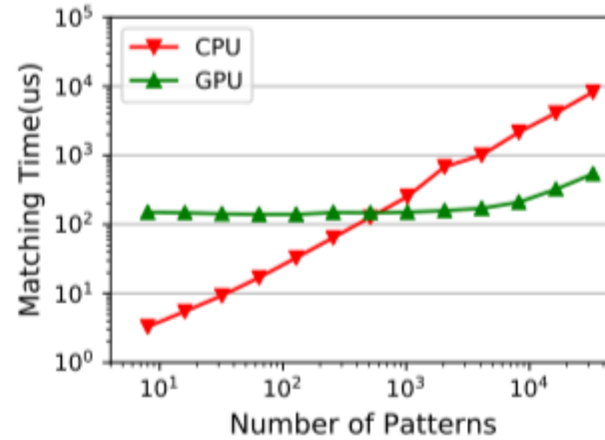
Integration

GPU

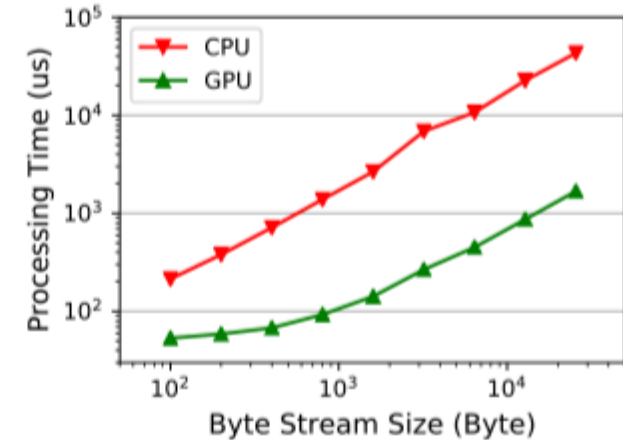
- + parallelism
- data transfer



(a) Encryption scaling up byte stream size



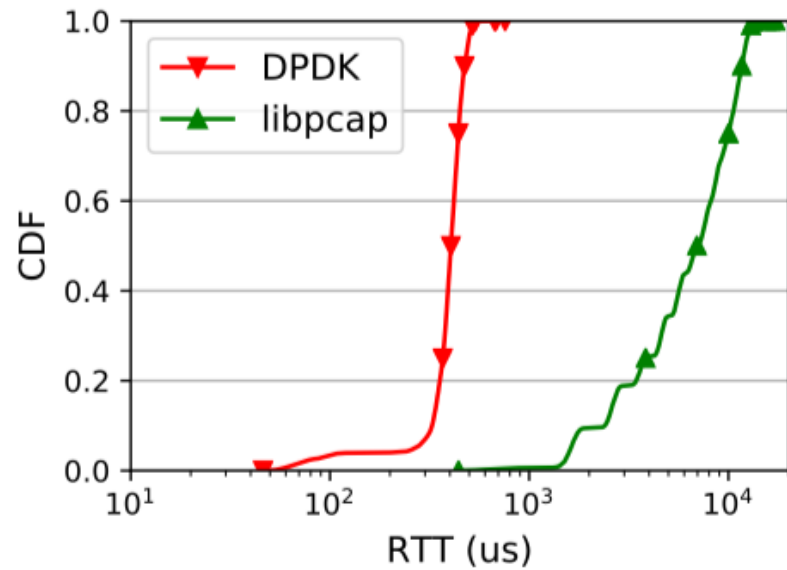
(b) Pattern matching scaling up number of patterns



(c) Pattern matching scaling up byte stream size

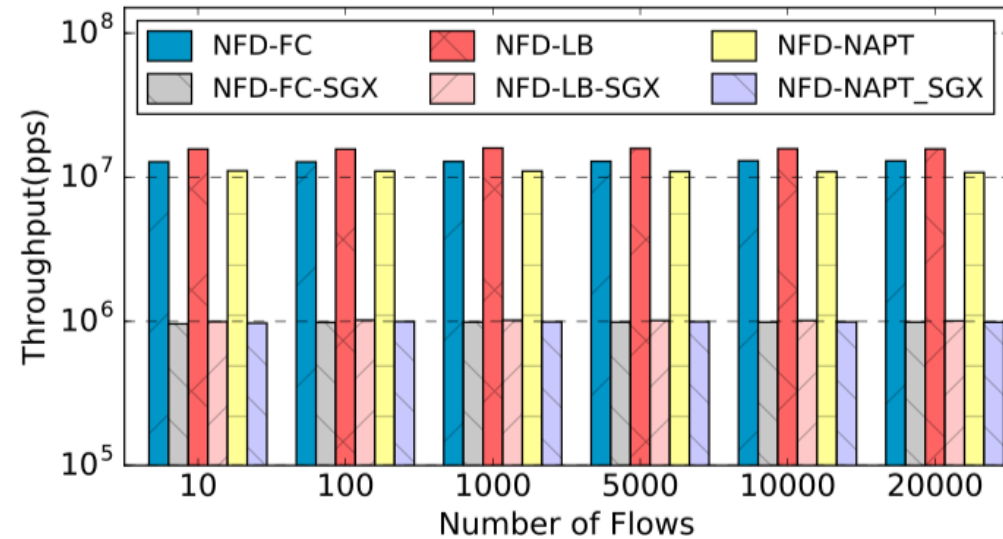
Integration

DPDK



405us v.s. 6952us

SGX



Content



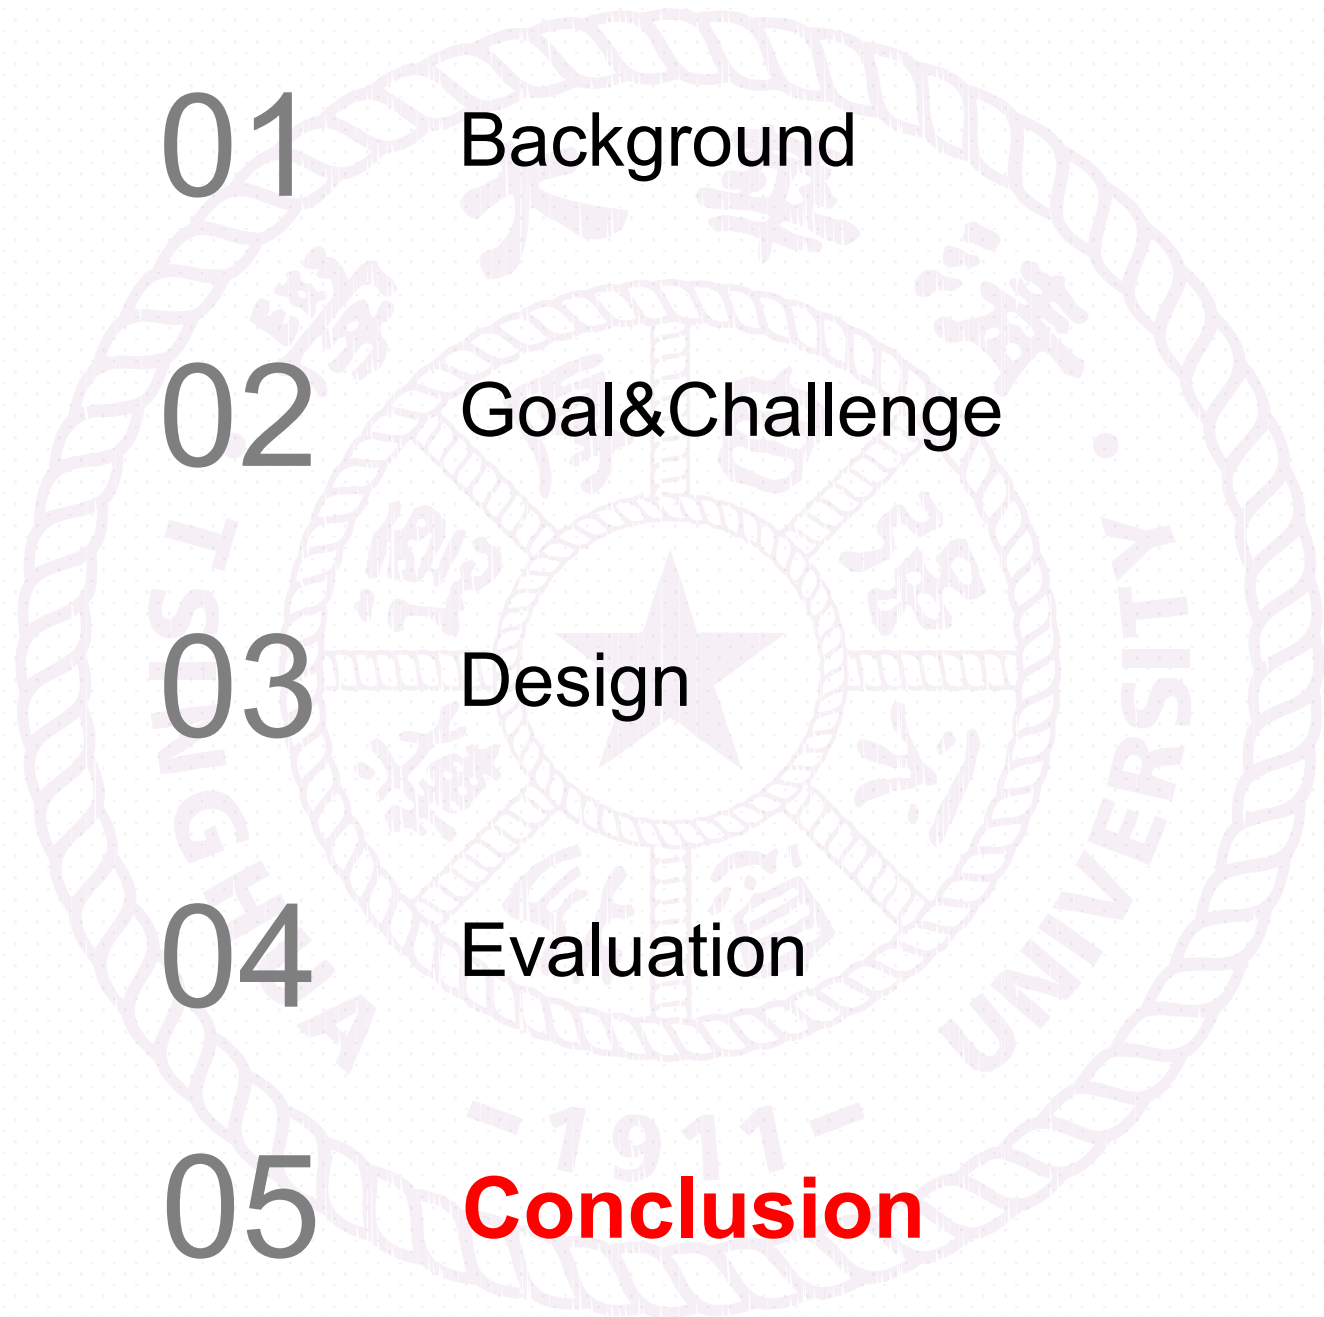
01 Background

02 Goal&Challenge

03 Design

04 Evaluation

05 **Conclusion**



Conclusion

We built a cross-platform NF development framework NFD

- Platform-independent language
- Reconfigurable compiler
- Develop 14 NFs with 6 platforms
- Less workload, valid logic and performance, platform compatibility, and commodity-equivalent complex logic



Thank You