



Microsoft Research

Checking Beliefs in Dynamic Networks


Nuno Lopes, Nikolaj Bjørner, Patrice Godefroid,
Karthick Jayaraman, George Varghese

Networks

- Business critical and complex

 Expensive bugs

- Fast protocol deployment in datacenters

 Frequent protocol changes

- A lot of legacy to maintain

 Operators don't have the full picture

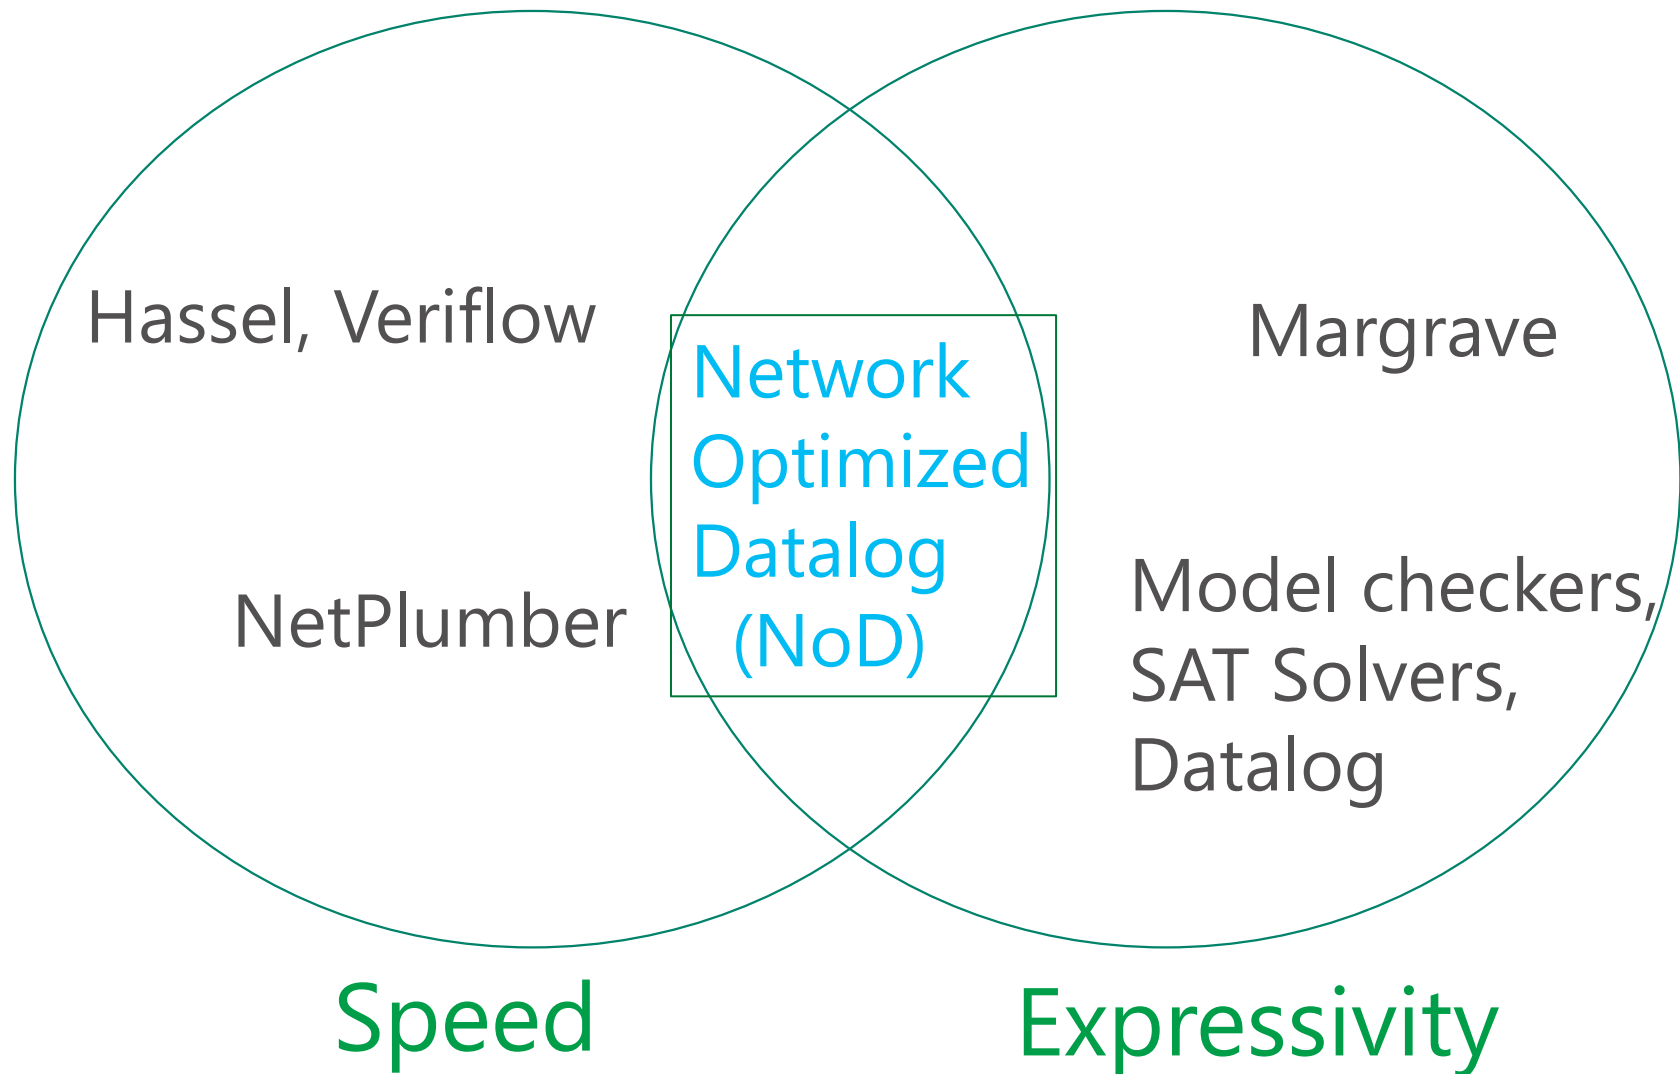
Network Verification to the Rescue

- Identify bugs
- Identify misbeliefs
- Increase confidence

This Paper in Context

- Implementation bugs (PIC)
 - E.g., protocol conformance
- Routing configuration errors (Batfish)
 - E.g., router filter error
- Dataplane configuration errors (NoD)
 - E.g., customer VMs can access controller

Existing Work versus Ours



Why Expressiveness Matters

- Network level
 - Enables modeling dynamic network behaviors such as new packet headers, new forwarding behaviors, failures, e.g.,
 - A P4 router adds a new header or a new forwarding behavior
- Specification level
 - Enables higher-level verification queries, e.g.
 - Customer VMs cannot reach fabric controller
 - All backup routers are equivalent

Example Beliefs

Policy Template	Example
Protection Sets	Customer VMs cannot access controllers
Reachable Sets	Customer VMs can access other VMs
Consistency	ECMP/Backup routes should have identical reachability
Middlebox	Forward path connections through middlebox should reverse
Locality	Packets between two hosts in the same cluster should stay within the cluster



Solution

Network-Optimized Datalog (NoD)

- Datalog for the specification of:
 - Data-plane/control-plane
 - Verification properties
- Tool for efficient verification
 - Available in open-source SMT solver Z3

Why Datalog?

- Good expressiveness/efficiency tradeoff
- Supports packet rewriting, load balancing
- Provides all (symbolic) solutions for “free”
 - Unlike SAT solvers or model checkers

Modeling Networks using Datalog

- Each matching rule in the FIB and each ACL rule becomes a Datalog rule
- State is set of packets at each router
- Packets start at sources; Datalog runs to fixed-point -> packets at destinations

So what's wrong with Datalog?

- Out-of-the-box implementations are slow
 - They work with a packet a time
- Our contributions:
 - Symbolic representation (dealing with sets of packets)
 - Efficient propagation of packets across routers

Symbolic Representation

- Packets represented as Difference of Cubes [NSDI'12]
- Generalized to support negation, useful e.g. to check consistency across backup routers

$$\bigcup_i \left(v_i \setminus \bigcup_j v_j \right)$$

v_i, v_j : ternary bit-vectors

Examples: $10\star 01\star \cup (10\star\star\star\star \setminus (10\star 01\star \cup \star\star\star 1\star\star))$
 $10\star 0\star\star$

Fuse Internal Datalog Operators

Packets with
Source starting
with 1



?

Evaluation


The background features a network of dark blue circles of various sizes connected by thin lines, set against a lighter blue gradient. A prominent large circle is in the upper right, with several smaller circles branching out from it and other circles scattered throughout the lower half of the frame.

Evaluation questions

- Do beliefs help?
- How hard is it to add a new forwarding protocol?
- How does NoD performs compared with existing verification tools?
- Is this useful in practice?

Beyond Reachability: Locality

- Found multiple violations of traffic locality



Query	Cluster 1	Cluster 2	Cluster 3
C2C	12 (2)	13 (2)	11 (2)
B2DSP	11 (2)	11 (2)	11 (2)
B $\bar{2}$ DSP	3 (1)	4 (1)	4 (1)
B2CSP	11 (2)	11 (2)	11 (2)
B $\bar{2}$ CSP	11 (2)	12 (2)	11 (2)

Verification time in seconds

Checking Operators' Beliefs

- Operators cannot specify reachability at VM level for millions of VMs
- They have “beliefs” of which sets of stations can reach others
- Found exceptions to operator’s beliefs
 - Customer VMs cannot access fabric controllers
- Process of belief refinement helps elicit specifications

Dynamism Example

- Experimental MPLS-like backbone with custom forwarding
- Took a few hours to model without any tool change
- Loop detection in < 1 second
- Identified 56 flows as black holes in 5 seconds

Performance Comparison

	Model Checker	SMT All Solutions	NoD	HSA
Stanford Unreach	12.2	0.1	2.1	0.1
Stanford Reachable	13.7	1121	5.9	0.9
Stanford Loop	11.7	290	3.9	0.2
Cloud	Time out	Time out	15.7	-
Cloud 2	8.5	Time out	4.8	-

Run time in seconds

Network Verification in Production

- Simplified version of NoD: SecGuru
 - Local checks on each router
- Deployed in Azure
- Finds ~1 problem per day

- Reduced legacy corporate ACL from 3,000 to 1,000 rules without outages

Conclusion

- NoD is expressive; takes as input:
 - Protocol specification -> Dynamism
 - Verification properties -> Beliefs
- More expressive than previous network verification tools, while competitive in speed
- Network operators' beliefs are fragile
- Code and benchmarks available on-line!

Microsoft Research