Deep Packet Inspection in FPGAs via Approximate Nondeterministic Automata

Milan Češka, Vojtěch Havlena, Lukáš Holík, Jan Kořenek, Ondřej Lengál, Denis Matoušek, **Jiří Matoušek**, Jakub Semrič, Tomáš Vojnar

Brno University of Technology, Faculty of Information Technology, IT4I Centre of Excellence, Czech Republic

Detection of Network Attacks

- The number of network attacks continuously increases.
- Detection of attacks is often performed by an NIDS (network intrusion detection system).
  - Monitors network traffic flowing through a given link.
  - Looks for characteristic patterns of known attacks.
  - Examples: Snort, Bro, Suricata.
- Attack patterns are usually described by a set of REs (regular expressions).
- Implementation of matching over an RE set typically utilizes a corresponding FA (finite automaton).
  - DFA (deterministic finite automaton).
  - NFA (nondeterministic finite automaton).

Requirements on NIDSes

- RE matching is the most demanding operation performed by an NIDS.
- Increasing trasfer rates require faster RE matching.
  - 100 Gbps equals to 6.72 ns per packet (worst case).
  - 400 Gbps equals to 1.68 ns per packet (worst case).
- Growing number of attack types require larger RE sets.
- To meet requirements on the performance of NIDSes, they have to use hardware-accelerated RE matching.

GPU-Based

- Utilizes either DFA or NFA.
- Theoretically sufficient performance for 100 Gbps links.
- Practically significantly limited by in/out throughput.
- High power consumption and latency.

FPGA-Based

- DFAs mapped into memory, NFAs mapped into logic.
- Various approaches for increasing throughput:
  - spatial stacking,
  - multi-striding.
- Approach based on parallel pipelined automata allows to scale throughput to 100 Gbps and beyond.

Parallel Pipelined Automata[1]

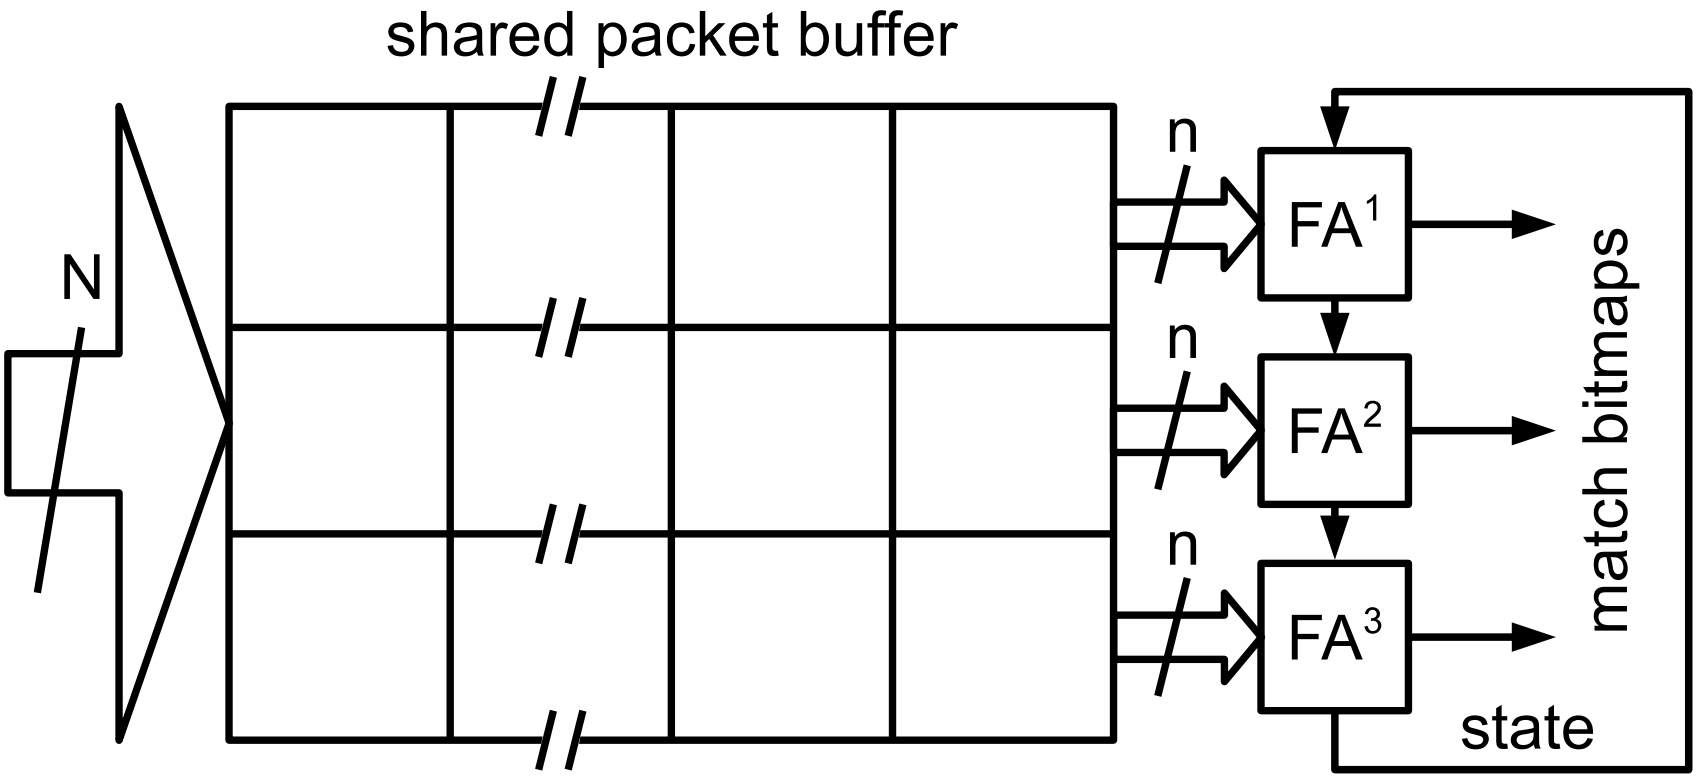

[1] D. Matoušek, J. Kořenek, and V. Puš, "High-speed Regular Expression Matching with Pipelined Automata," in FPT'16. IEEE, 2016.

- Implementations for both DFAs[2] and NFAs[1] exist.
- Throughput and resources scale linearly with the number of pipeline stages.
  - 100 Gbps requires 64 stages (8 b @ 200 MHz).
  - 400 Gbps requires 256 stages (8 b @ 200 MHz).
- Allows to implement RE matching over a simple RE set with 100 Gbps throughput in a single FPGA chip.
- Problematic for large RE sets.
  - Each stage contains an FA representing a full RE set.

[2] D. Matoušek, J. Kubiš, J. Matoušek, J. Kořenek, "Regular Expression Matching with Pipelined Delayed Input DFAs for High-speed Networks," in ANCS'18. ACM, 2018.

Simulation-Based

- A well-studied approach with several variants.
- Preserves the original automaton's language.
- Usually only a limited reduction of automaton's states and transitions.

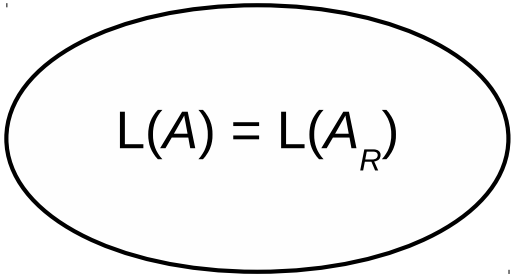

Approximation-Based

- A modern approach with ongoing research.
- Over-approximates the original automaton's language.
  - Produces false positives.
- Variable reduction of automaton's states and transitions.

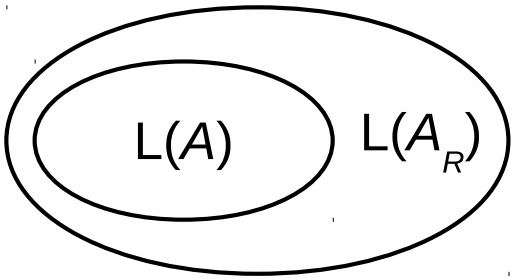

Multi-Stage Architecture

- Several stages of gradually more precise NFAs.
  - Each stage is a pipeline of parallel automata.
- A larger number of less precise NFAs pre-filter traffic for a lower number of more precise NFAs.

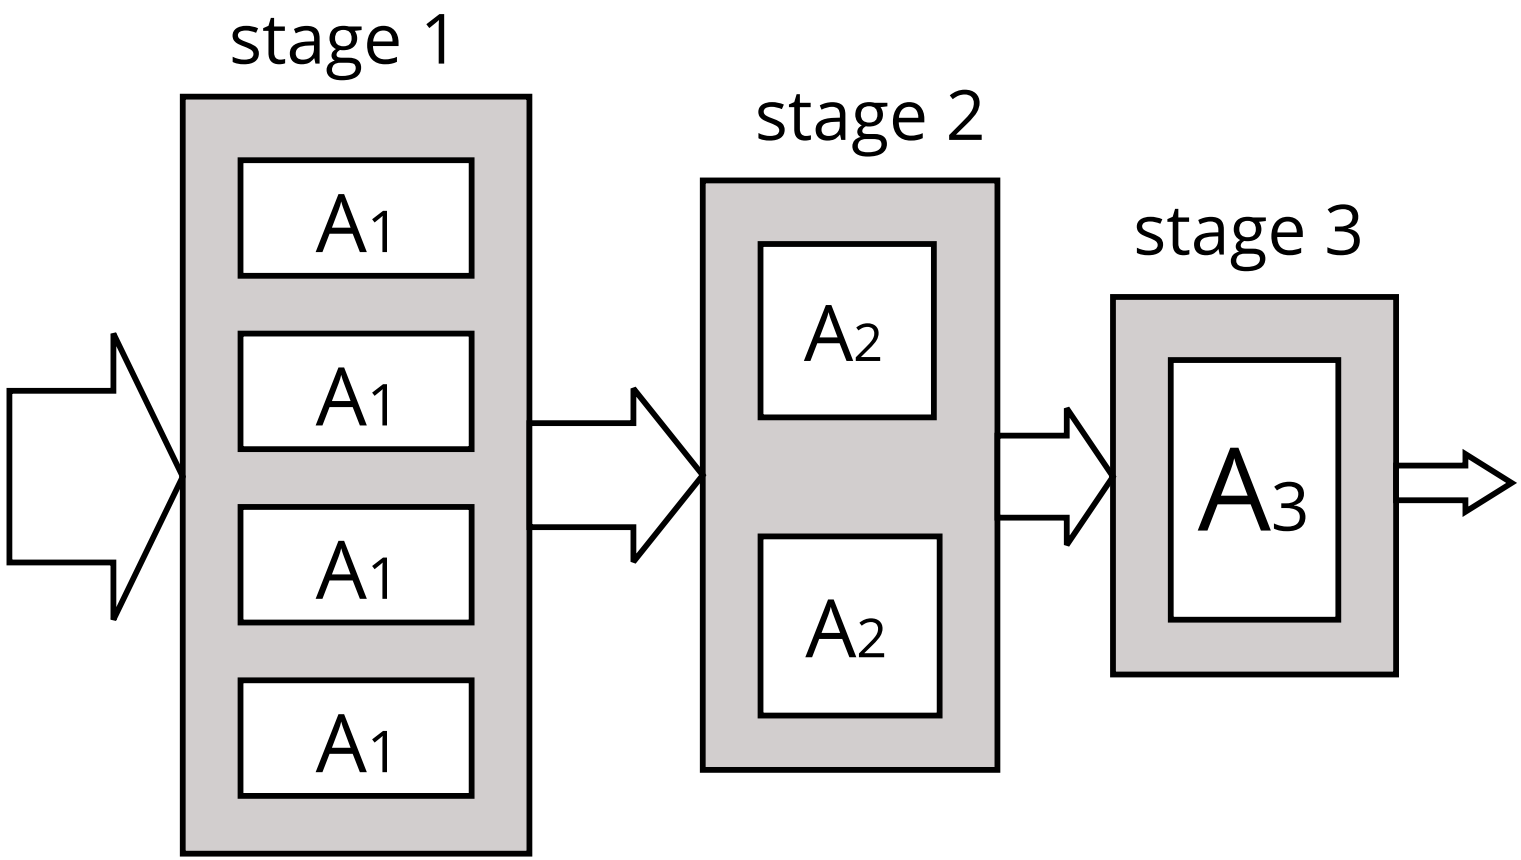

- 100 Gbps input (512 b @ 200 MHz).
- Three approximated NFAs with 6.4 Gbps throughput (32 b @ 200 MHz).
  - $\mathcal{A}_3$ is the precise NFA $\mathcal{A}$.

	LUT	*Acc*
$\mathcal{A}_1$	100	0.5
$\mathcal{A}_2$	200	0.2
$\mathcal{A}_3$	1,000	0.1

- Possible configurations of the architecture
  - 10,000 LUTs are available.

#	Stg. 1	Stg. 2	Stg. 3	LUTs	output
1	$16\times\mathcal{A}_3$	—	—	16,000	10 Gbps
2	$16\times\mathcal{A}_2$	$4\times\mathcal{A}_3$	—	7,200	10 Gbps
3	$16\times\mathcal{A}_1$	$8\times\mathcal{A}_3$	—	9,600	10 Gbps
4	$16\times\mathcal{A}_1$	$8\times\mathcal{A}_2$	$4\times\mathcal{A}_3$	7,200	10 Gbps

Pruning reduction

- Prunes out insignificant (rarely visited) states.

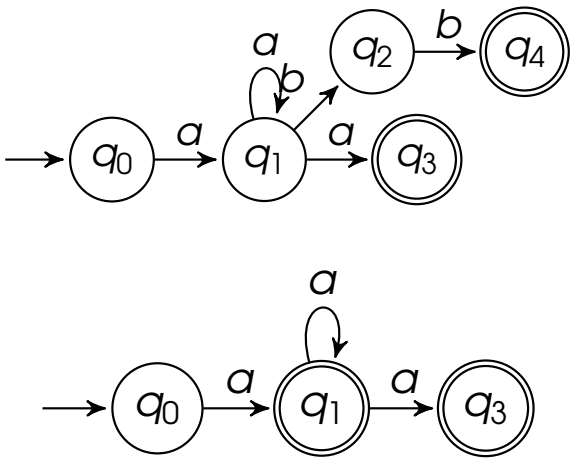

Merging reduction

- Merges similarly significant adjacent states.

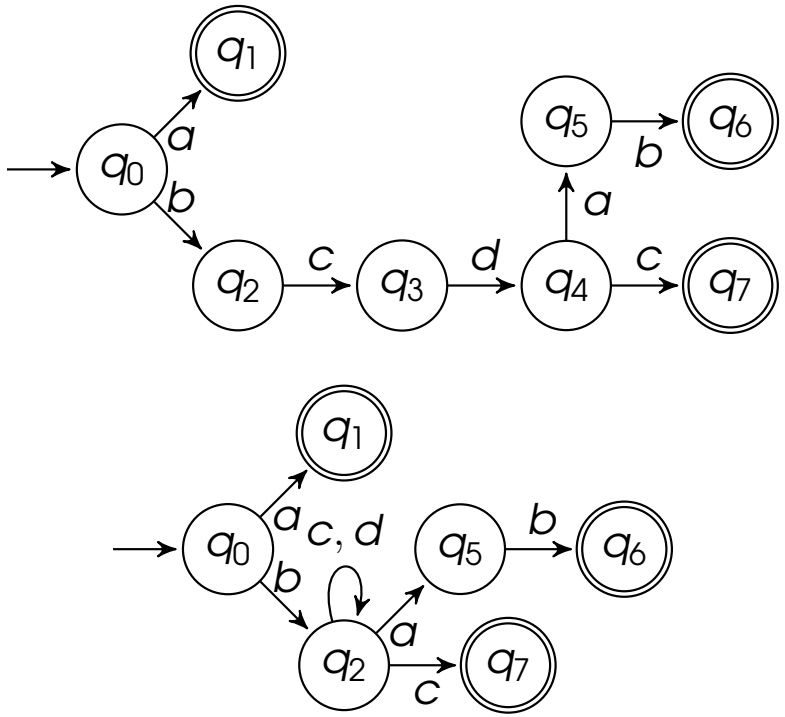

- Reductions can be controlled via a reduction ratio.

Acceptance precision

$$AP = \frac{A_{TP}}{A_{FP}+A_{TP}}$$

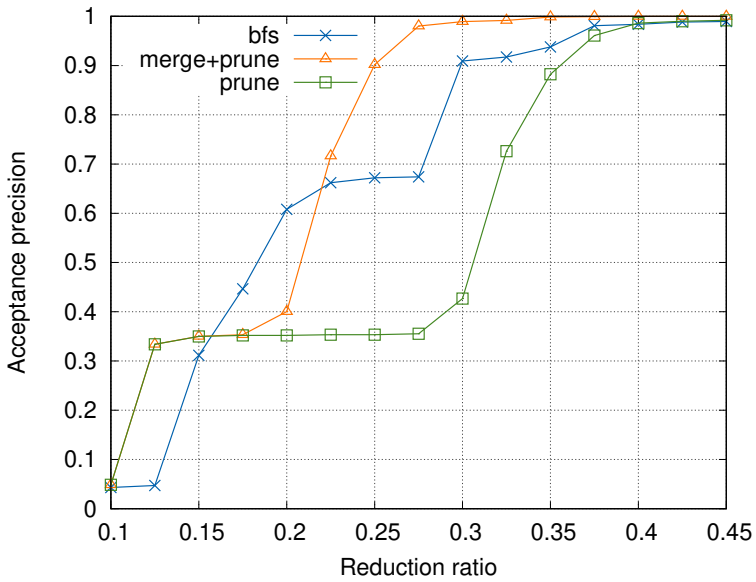

Acceptance probability

$$Prob = \frac{A_{TP}+A_{FP}}{|\mathcal{S}|}$$

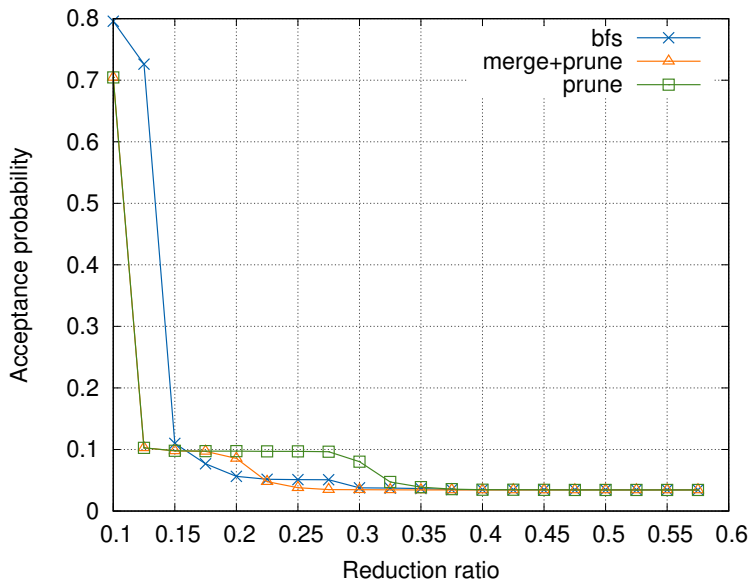

`spyware` RE set (461 REs; 12,809 states; 279,334 transitions)

A_{TP} . . . acceptance true positives
A_{FP} . . . acceptance false positives
$|\mathcal{S}|$ . . . the size of network traffic sample

- Target FPGA: Xilinx Virtex UltraScale+ VU9P.
  - 1,182k LUTs in total.
  - 737k LUTs available for RE matching.

Precise				
speed	1 stg	2 stg	3 stg	4 stg
100	5M	444k	296k	296k
200	10M	809k	513k	513k
400	20M	1.5M	945k	945k

4 % of traffic				
speed	1 stg	2 stg	3 stg	4 stg
100	227k	61k	65k	69k
200	453k	122k	126k	133k
400	907k	242k	247k	261k

`spyware` RE set (461 REs; 12,809 states; 279,334 transitions)

Summary of Contributions

- Approximation-based techniques for reduction of NFA's states and transitions.
- A multi-stage architecture implementing RE matching based on approximated NFAs organized in processing pipelines.
- Combination of the previous contributions in a hardware-accelerated RE matching system supporting large RE sets and throughput beyond 100 Gbps implemented on a single Virtex UltraScale+ FPGA.

Thank you for your attention.