

Xiaodong Yu¹, Michela Becchi² (Master's advisor), and Danfeng (Daphne) Yao¹ (PhD advisor)

1. Dept. of Computer Science, Virginia Tech, USA {xdyu, danfeng}@vt.edu 2. Dept. of ECE, North Carolina State University, USA mbecchi@ncsu.edu

Introduction and Problem Statement

- Automata-based processing has been applied to diverse areas as the core of the searching engine:
 - network security, NLP, data mining, bioinformatics.
- There are two primary representations of automata:
 - NFA: has excellent spatial efficiency, but may require expensive per-character processing.
 - DFA: offers limited per-character processing at the cost of a large automaton (state explosion (SE) issue).
- Current SE issue solutions suffer from some limitations:
 - only work on small and simple datasets.
 - have prohibitive worst-case processing time.
 - can't guarantee fully functional equivalence.
- The searching engine of NIDS using conventional NFA/DFA can't directly handle the out-of-order packets.
 - it has to buffer the packets and then process flow only after reordering and reassembling the packets.
 - this is vulnerable to denial-of-service (DoS) attacks.
- Automata processing is inherently difficult to be parallelized especially on the many-core (e.g., GPU) devices due to:
 - tight dependencies between every two computation steps.
 - unstructured memory-access pattern.
- The newly proposed reconfigurable device – Micron's Automata Processor (AP) suffers from two problems:
 - Current APIs require the domain end-users to have insights of both AP architecture and automata processing.
 - Requires vast reconfiguring cost for a large dataset.
- It's unclear the advantages and disadvantages of different automata processing accelerators and the innovation space
 - CPU, GPU, FPGA, Automata Processor, etc

Contribution Summary

- Algorithm Level:**
 - A1. JFA: an FA variant that uses state variables to avoid SE, and is functionally equivalent to the corresponding DFA [2]
 - A2. O³FA: an automata-based DPI engine for RegEx matching on out-of-order packets without requiring flow reassembly [3]
- Implementation Level (frameworks):**
 - F1. A GPU-based automata processing framework that supports different automata representations and provides various architecture-aware optimization techniques [1]
 - F2. Robotomata: a framework for auto APM-automata codes generation and optimization and fast reconfiguration on AP [4]
 - F3. A toolchain that allows the apple-to-apple comparison of NFA acceleration engines on different platforms [5]

Background Knowledge

Finite Automata

RegEx: (1) $a+bc$ (2) $bcd+$ (3) cde

NFA

DFA

State Explosion

NFA: N states → DFA: 2^N states

In practice, a desktop machine equipped with 2GB of memory will report an out-of-memory error after generating **1,000,000 DFA states**, which would be expected for a typical **100-state NFA**.

Automata Processor

Architecture Overview

Programming Interface

ANML: Automata Network Markup Language

- has **programmability issue**
- very low-level and complex, requires programmers to manipulate STEs

Workflow

DRAM-based Design

States encoded in State Transition Elements (STEs) = 256-bit masks

Programmable Resources: 49,152 STEs = one core (chip) → 1.5 MB/chip

16, 32 or 48 cores/board

STE: 49152 per chip

Counter Element: 768 per chip

Boolean Element: 2304 per chip

STATES STORED IN MEMORY

TRANSITIONS STORED IN ROUTING MATRIX

A1. JFA

Basic JFA Design

Optimization 1: Label Forwarding: aim to limit the label propagation to reduce the number of state and condition complexity

Optimization 2: Label Splitting: aim to reduce the label splitting condition complexity when label forwarding is not possible

JFA with Label Splitting

Experiment Results

The characteristics of JFA compared to NFA, DFA, Hybrid-FA (for Snort datasets)

Dataset	NFA		Small DFA (<100k states)		Large DFA		Hybrid-FA		JFA							
	# states	trans.	# states	trans.	# states	trans.	# states	trans.	# states	trans.						
Backdoor	2,968	-41k	9	-148k	7	-172k	-1,527k	3,731	924	72	317	126,917	5,615	-2,008k		
Spyware	4,690	-53k	11	-366k	7	-2,024k	-3,880k	-15,200k	1,636	3,690	45	56	18,867	1,358	-71k	
Dostar.05	14,978	-70k	8	-268k	5	-1,596k	5	-604k	-3,421k	75,540	5,022	105	2,116	36,537	162	-122k
Dostar.1	14,688	-79k	10	-677k	7	-3,414k	7	-2,346k	-10,386k	65,536	6,187	116	2,700	35,876	384	-121k
Dostar.2	14,249	-88k	21	-736k	12	-3,498k	12	-6,808k	-31,471k	27,904	6,965	149	1,995	34,762	1,616	-123k

A2. O³FA

Out-of-Order Packets issue in NIDS

In Network Intrusion Detection Systems (NIDS), input streams have size at MB/GB level, whereas single packet payload is only 64KB

- divided to multiple substreams → carried by different packets

Out-of-Order Packets issue: packets arrive in a random order, so match can escape from the NIDS due to cross-packet matching

input stream: $abcde fgh$ malicious pattern: $cdef$

Normal order: $abcde fgh \rightarrow abcde fgh$ match $cdef$

Out of order: $abcde fgh \rightarrow efgh \rightarrow abcd$ no match

Traditional Solution

Buffering and Reassembling

Drawbacks:

- Very resource-intensive
- Attackers can exhaust the packet buffer capacity to cause DoS

O³FA Design

Regular-DFA: Prefix-NFA, Non-anchor Suffix-NFA

Regular-FA: Anchor Suffix-NFA

Evaluations

Buffer size savings:

- O³FA states buffer vs. traditional packets buffer
- State buffer requires 20x-4000x less size

Bandwidth overhead:

- traversed NFA states #
- traversed state # in reassembly methods

Traversal overhead:

- O³FA traversal overhead
- traversed state # in reassembly methods

F1. GPU-based Automata Processing

General GPU-based Design

GPU-based NFA

Baseline design

Optimization #1: Sort transition vector by source state

Optimization #2: Partition states into compatibility groups

GPU-based DFA

RegEx Partitioning

Compressed DFA

Software-managed Cache

Using shared memory as cache to better leverage the temporal locality of state traversal

Experiment Results

average traffic

malicious traffic

F2. ROBOTOMATA (Framework for Automata Processor)

Paradigm-based Hierarchical Construction

Aim to mitigate the programmability issue

The above design is in back end and transparent to end users!

Cascadable Macros

Aim to mitigate the scalability issue

The above design is in back end and transparent to end users!

End-users only need to provide: (error types, pattern length, error)

End-users only need to: Enable cascadable macros

Evaluations

ROBOTOMATA can achieve up to **46.5x** compilation time speedups over the ANML

F3. Platform Comparison Toolchain

Goal of the toolchain

Apple-to-apple comparison between automata processing accelerators

- 3 platforms: GPU, AP, FPGA
- non-deterministic finite automata (NFA)
- Large datasets (million states)

Dimensions of Comparison

- Resource requirements
- Traversal throughput
- Preprocessing time

Toolchain Design

Platform Summary

	GPU	Micron's AP	FPGA
ENCODING/states	memory	memory	logic
ENCODING/transitions	memory configuration	memory configuration	logic
PROGRAMMING	strongly dataset dependent	weakly dataset dependent	compilation, place&route
TRAVERSAL BEHAVIOR	strongly dataset dependent	weakly dataset dependent	weakly dataset dependent

Evaluations

Resource utilization

Traversal throughput

Preprocessing time

FPGA: best traversal throughputs; significant preprocessing (aggregate GPU: modest (aggregate) traversal throughputs; low preprocessing time; good pattern densities

AP: intermediate btw FPGAs and GPUs; most suited many small NFAs with fixed topology

References

- X. Yu and M. Becchi, "GPU Acceleration of Regular Expression Matching for Large Datasets: Exploring the Implementation Space," In ACM CF'13
- X. Yu, B. Lin and M. Becchi, "Revisiting State Blow-up: Automatically Building Augmented-FA while Preserving Functional Equivalence," Selected Areas in Communications, IEEE Journal on (JSAC.2014)
- X. Yu, W.-c. Feng, D. Yao and M. Becchi, "O³FA: A Scalable Finite Automata-based Pattern-Matching Engine for Out-of-Order Deep Packet Inspection," In ACM/IEEE ANCS '11
- X. Yu, H. Wang, K. Hou, W.-c. Feng, "ROBOTOMATA: A Framework for Approximate Pattern Matching of Big Data on an Automata Processor," In IEEE BigData'17
- M. Nourian, X. Wang, X. Yu, W.-c. Feng and M. Becchi, "Demystifying automata processing: GPUs, FPGAs or Micron's AP?," In ACM ICS '17