Hardware Acceleration of Algorithms in Computer Networks Using FPGA

Jan Kořenek

Brno University of Technology, Faculty of Information Technology
Božetěchova 2, 612 00 Brno, CZ
www.fit.vutbr.cz/~korenek

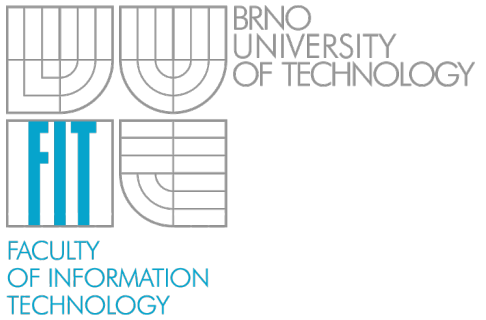

INVESTMENTS IN EDUCATION DEVELOPMENT

Outline

- Introduction and Motivation
- Hardware acceleration of time critical operations
- HW framework for research and development
- Applications for ISP, companies and LEA
- Conclusions and Future work

- Wire speed processing is needed in many network applications:
  - Low performance can be utilized by intruder to hide an attack to the Intrusion Detection System (IDS)
  - NetFlow statistics with packet loss or packet level sampling significantly decrease precision of network behavioural analysis
  - Lawful interception system has to capture all requested packets without any packet loss even for heavy loaded links
  - ....

Example: Network Intrusion Detection System

Data stream is transferred by many packets

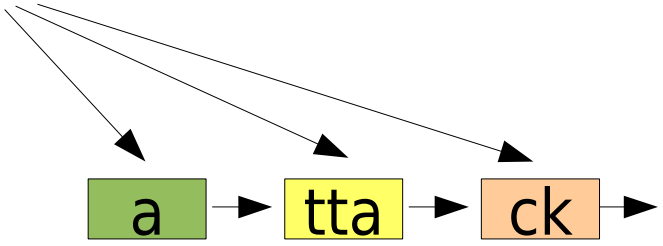

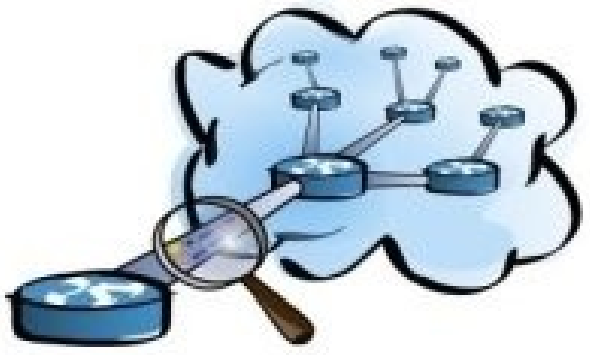

Signature "attack" is used to detect threat

If packet is dropped IDS miss detection

- Time to process one packet

Line rate [b/s]	1G	10G	40G	100G
Maximal Packet rate [Packets/s]	1,5M	15M	60M	150M
CPU 3GHz [clock cycles/packet]	1807	181	45	18

181 clock cycles per packet even for 10G links

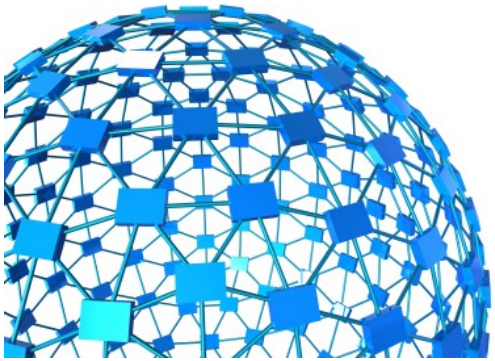

- Processor performance for time-critical operations

Operation	Throughput	1G	10G	40G	100G
Protocol Parsing	21Mp/s	✓	✓	STOP	STOP
TCP Stream Reassembly	2,5Mp/s	✓	STOP	STOP	STOP
Packet Classification	12Mp/s	✓	STOP	STOP	STOP
Pattern Matching	400Mb/s	STOP	STOP	STOP	STOP

Results for one processor core

- ***We can split traffic among multiple cores (RSS,TSS)***

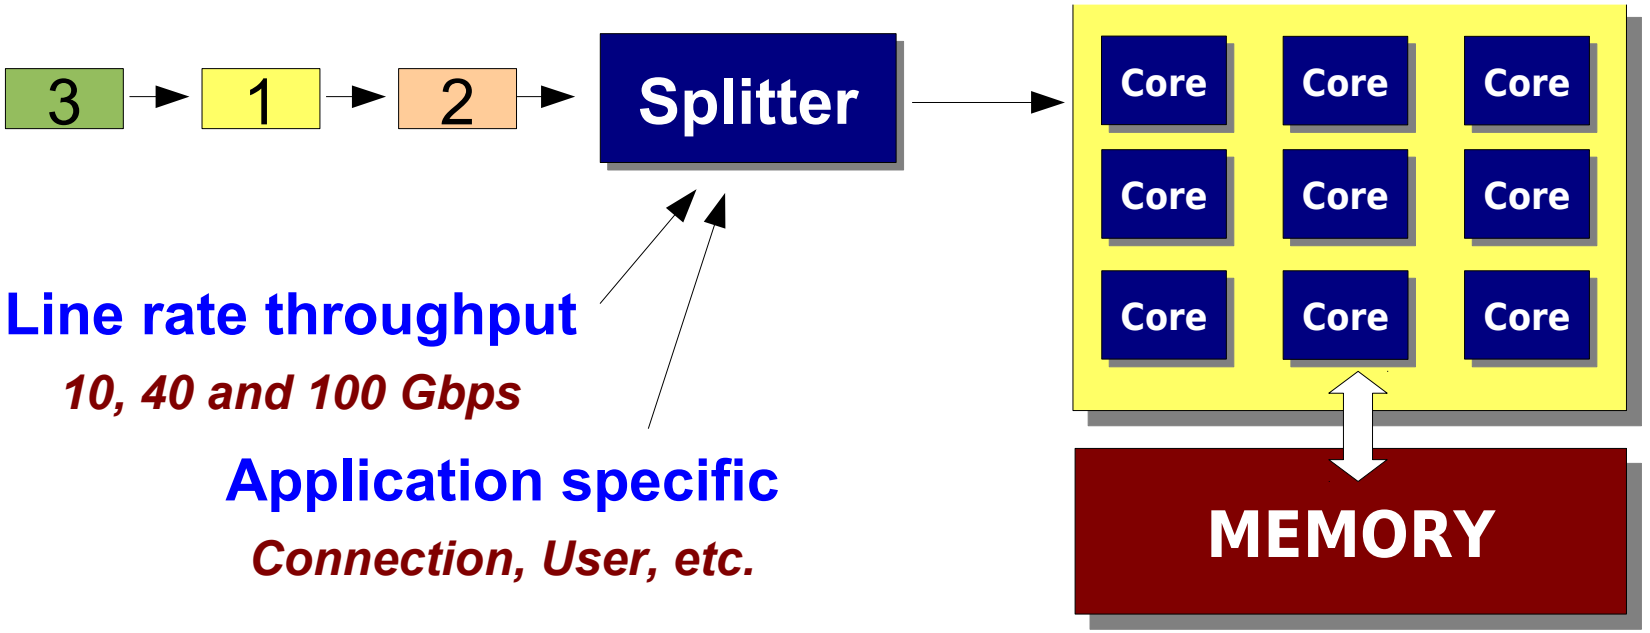

- ***Performance can not be multiplied by number of cores!***
  - Many applications can not be speed up by parallel processing
  - Communication overhead between cores

- Packet header analysis and header fields extraction
  - Parsing of packet headers and extraction of selected header fields
  - Any network device need packet header parsing (MAC or IP addresses)
- Longest prefix match or IP lookup
  - Find longest prefix for an IP address
  - Core routers have usually more than 300 thousand of IP prefixes in routing tables
- Packet classification
  - Find classification (filtering) rule for every received packet
  - Packet filters have usually hundreds or thousands rules related to multiple packet header fields (IP addresses, Ports, Protocol, etc.)
- Pattern matching
  - Matching strings or regular expressions in packet payload
  - IDS Snort have thousands of rules with strings and regular expressions

- Extraction of packet header fields for packet filter, etc.
- Packet Headers Analysis – identification of header fields
  - Hardware architecture generated from XML configuration
  - DFA modified to process multiple bytes per clock cycle
  - Configurable data width to address trade-off between hardware resources and processing speed
- Headers Field Extraction – extraction of selected data
  - Hardware architecture configurable at runtime by XML

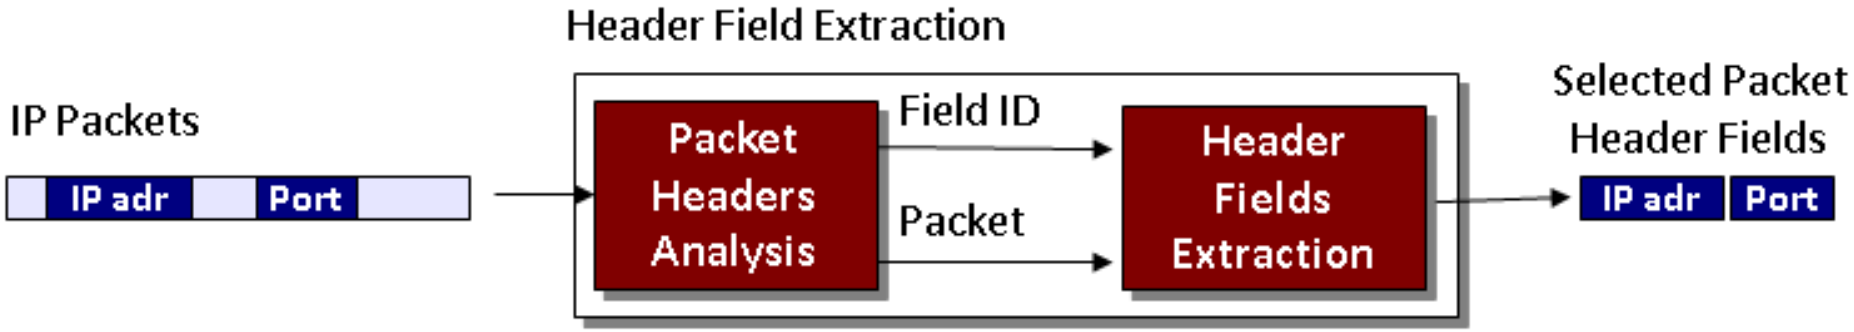

- Find longest prefix for IP address. Core routers have routing table with more than 300k prefixes

- Related Algoritms
  - Trie,TreeBitmap, Shape Shifting Trie, ...
- Prefix Representation
  - TreeBitmap – single node ~8B, up to 8 nodes for IPv4 representation (64B), 8 steps - engine replication required
  - Whole IPv4 lookup using Hash (4B per IP) – can not be applied to prefixes

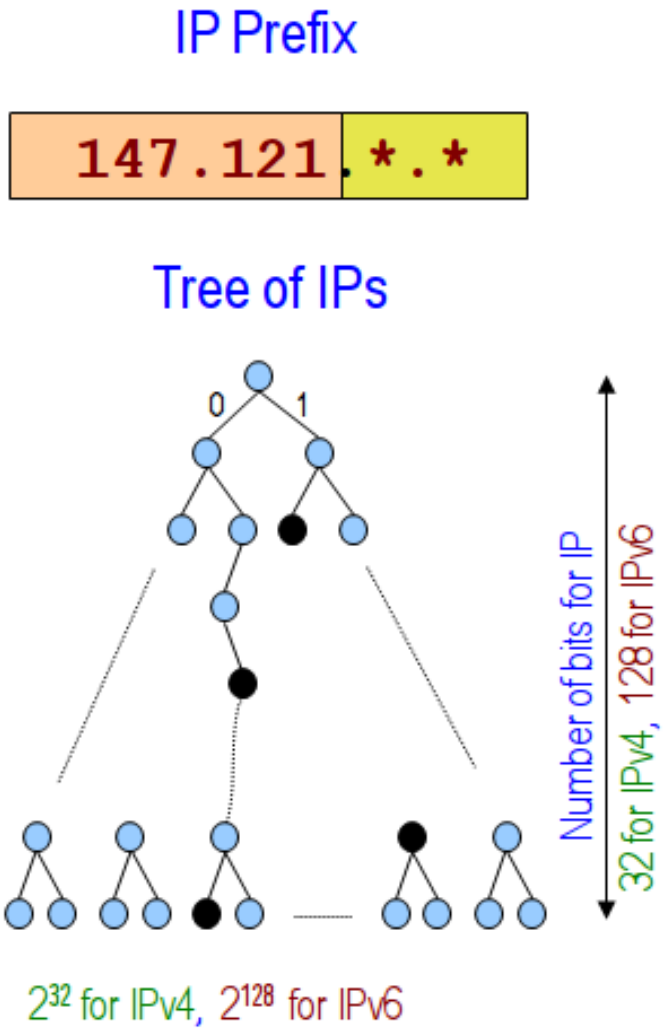

- Tree Bitmap only is unusable
  - 128/4 = 32 Steps
  - Massive unit replication
- Hash for unique IP address
  - 128 bits (16B) for IP
  - Lookup in single step
- Combination of Hash and Tree Bitmap

  100 Gbps throughput even for IPv6 networks

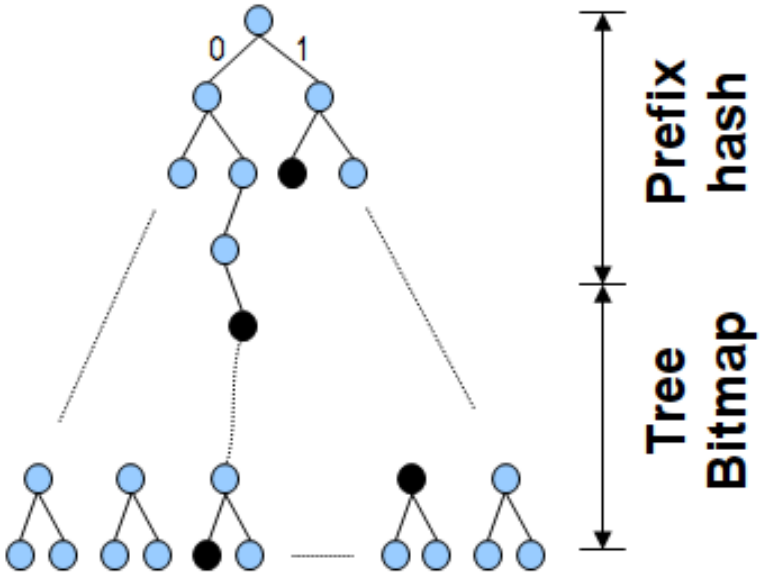

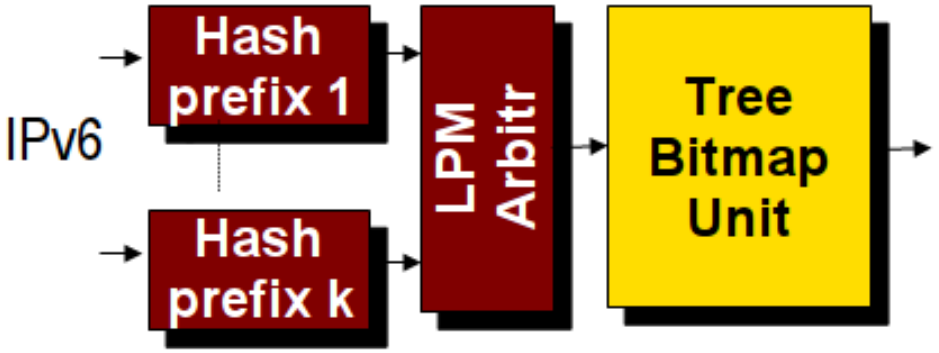

- Packet classification performs decision about each packet based on the values found in the packet header
- Packet header fields are *dimensions* in classification rules
- Condition types:
  - Exact match (Protocol)
  - Prefix (IP address)
  - Range (Ports)

Classification rules

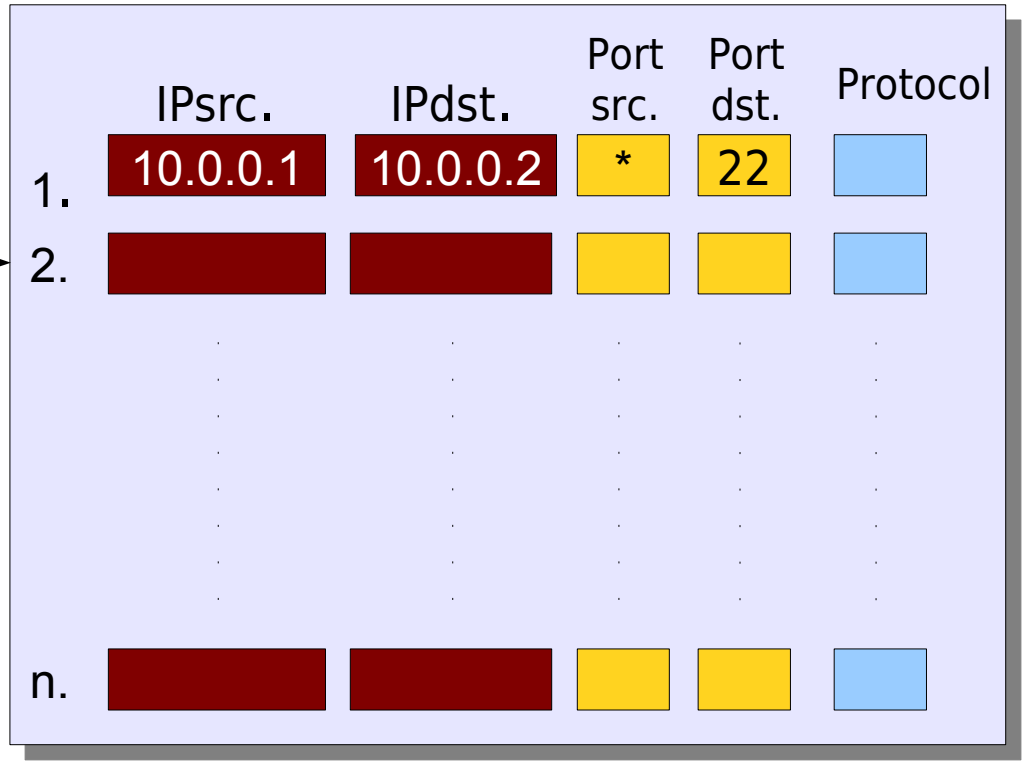

Header fields are dimensions

Received packet

Lookup time

- 10 Gbps ~ 50 ns
- 100 Gbps ~ 5 ns

- Perfect Hash Crossproduct Algorithm (PHCA)
  - Longest Prefix Match (LPM) on selected header fields
  - Search filtering rule by Perfect Hash
  - Comparison of matched rule to header fields
- Constant time complexity
- Reduction of memory requirements by intended collisions of perfect hash function

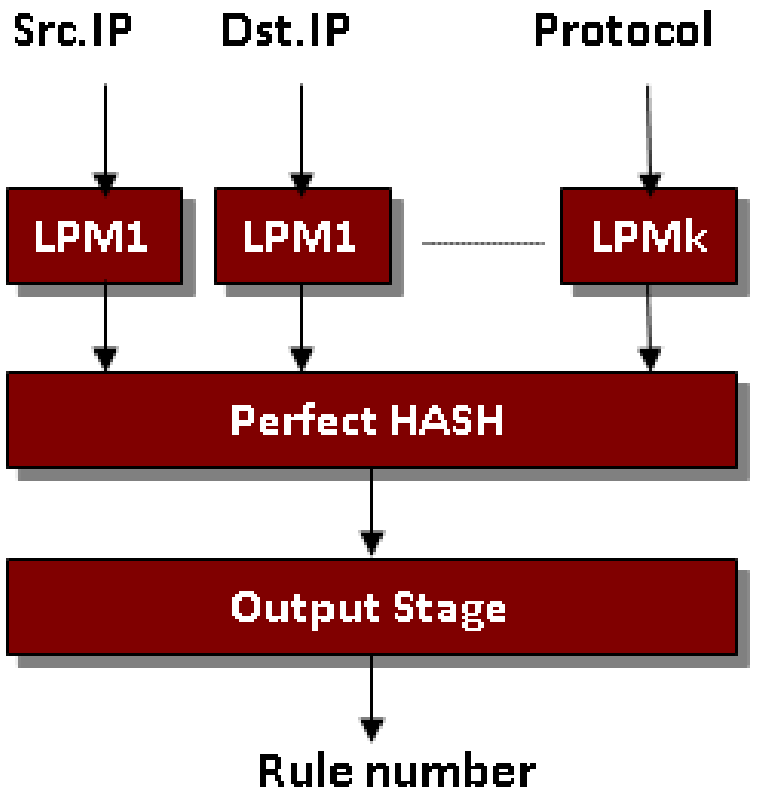

More than 100 Gbps throughput with reasonable memory requirements (2x QDR SRAM)

- IDS Snort have thousands of rules with strings and regular expressions
- Deterministic Finite Automaton (DFA) based architectures
  - Exponential growth of transition table → ***high memory requirements***
  - Partitioning regular expressions into multiple groups
  - Size of transition table is reduced by implicit transitions (Delayed Input DFA)
- Nondeterministic Finite Automaton (NFA) based architectures
  - Mapping of NFA to FPGA Logic → ***requirements for large FPGA capacity***
  - Shared decoder can significantly reduce FPGA logic utilization

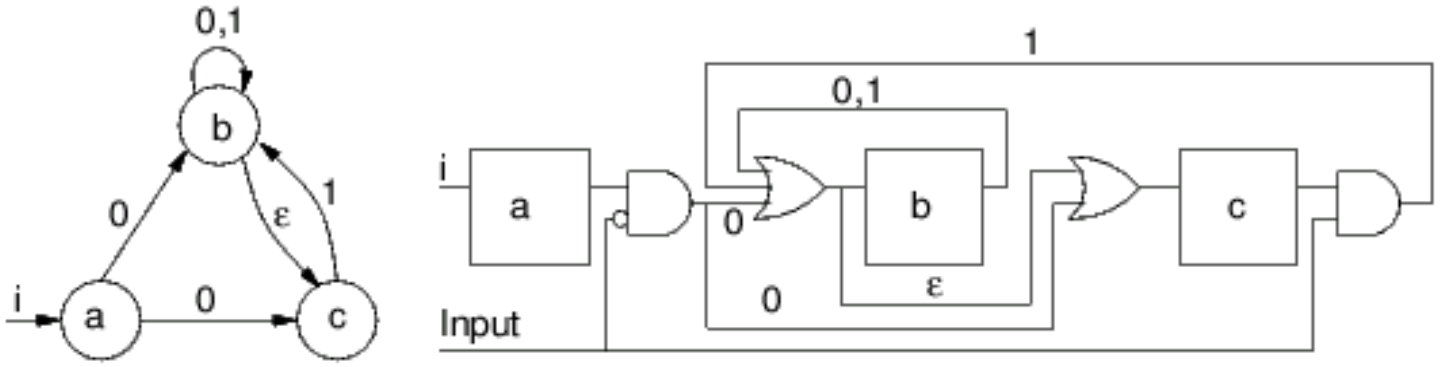

Example of NFA representation in a FPGA logic

- For IDS Snort only 3% of all states can be concurrently active
- ***Set of States without Collisions*** (Q^{NCA}) – two different states can not be active at the same time
- Set of states without collisions can be represented by DFA unit
  - Transitions are stored in memory instead of LUTs
  - Binary encoding of states significantly reduce amount of FF
- Other states and transitions are mapped to FPGA logic

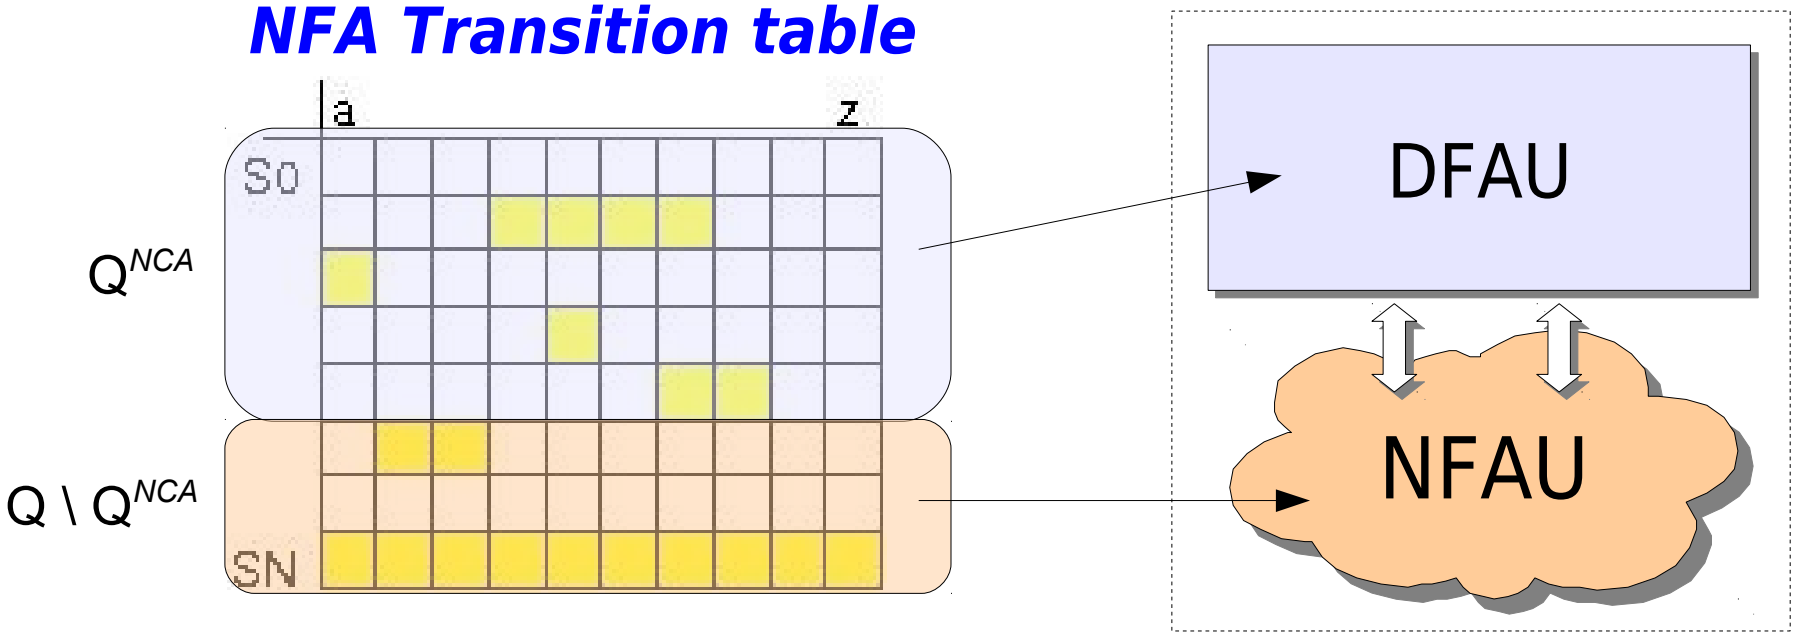

- FPGA Logic Utilization for Xilinx Virtex-5 LX155T

	Clark et al		NFA Split			Reduction	
	LUT[-]	FF[-]	LUT[-]	FF[-]	BRAM[-]	LUT[%]	FF[%]
L7 Dec.	1 538	836	1 231	237	3	80.04	28.35
Snort (1)	4 680	4 043	2 166	821	21	52.69	20.31
Snort (2)	2 965	876	1 883	374	15	63.51	42.69
Snort (3)	1 637	555	1 370	261	6	83.69	47.03
Snort (4)	2 436	1 392	2 233	924	6	91.67	66.38
Snort (5)	2 807	1 099	1 969	368	6	70.15	33.48
Snort (6)	2 680	1 097	2 259	543	6	84.29	49.50
Snort (7)	10 314	2 812	3 393	1 439	6	32.90	51.17
Snort (8)	18 565	13 042	11 332	5 484	30	77.20	42.05
Snort (9)	65 265	40 910	29 670	20 147	39	45.46	49.25
Snort (all)	111 349	65 826	59 575	30 361	135	53.50	46.12

NFA Split architecture reduces 66.8% LUT and 43.3% flip-flops in average for all sets of regular expressions

- NetCOPE – high performance scalable framework for rapid development of FPGA applications and rapid prototyping
- The framework was developed in cooperation with CESNET and is provided by INVEA-TECH company

- Basic Features
  - 1 Gbps and 10 Gbps Ethernet support
  - PCI Express x1, x4, x8
  - Hardware abstraction layer
  - LocaLink protocol for data flows

Supported Boards

COMOv2

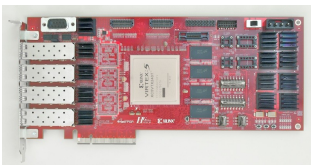

NetFPGA 10G

- Packets are stored in a ring buffer in RAM
- Data format is independent on the application
- Whole communication is driven only with buffer StartPointer, EndPointer
- Processing models
  - Every thread has to process all data in buffer
  - Distribution of data among multiple threads

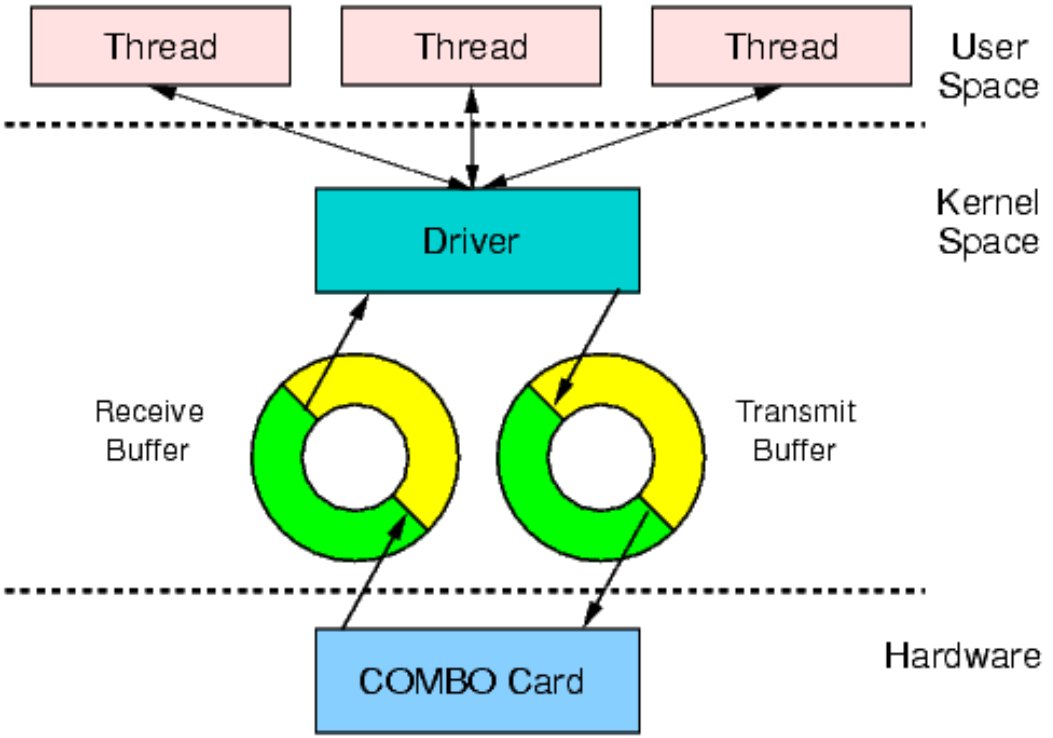

Fully saturated 10 Gbps link of network traffic can be transferred to the host memory without packet loss

- Hardware acceleration of NetFlow monitoring
  - Header fields extractions, pattern matching, hash computation and distribution among multiple cores are implemented in hardware
  - FlowCACHE management implemented on processor cores

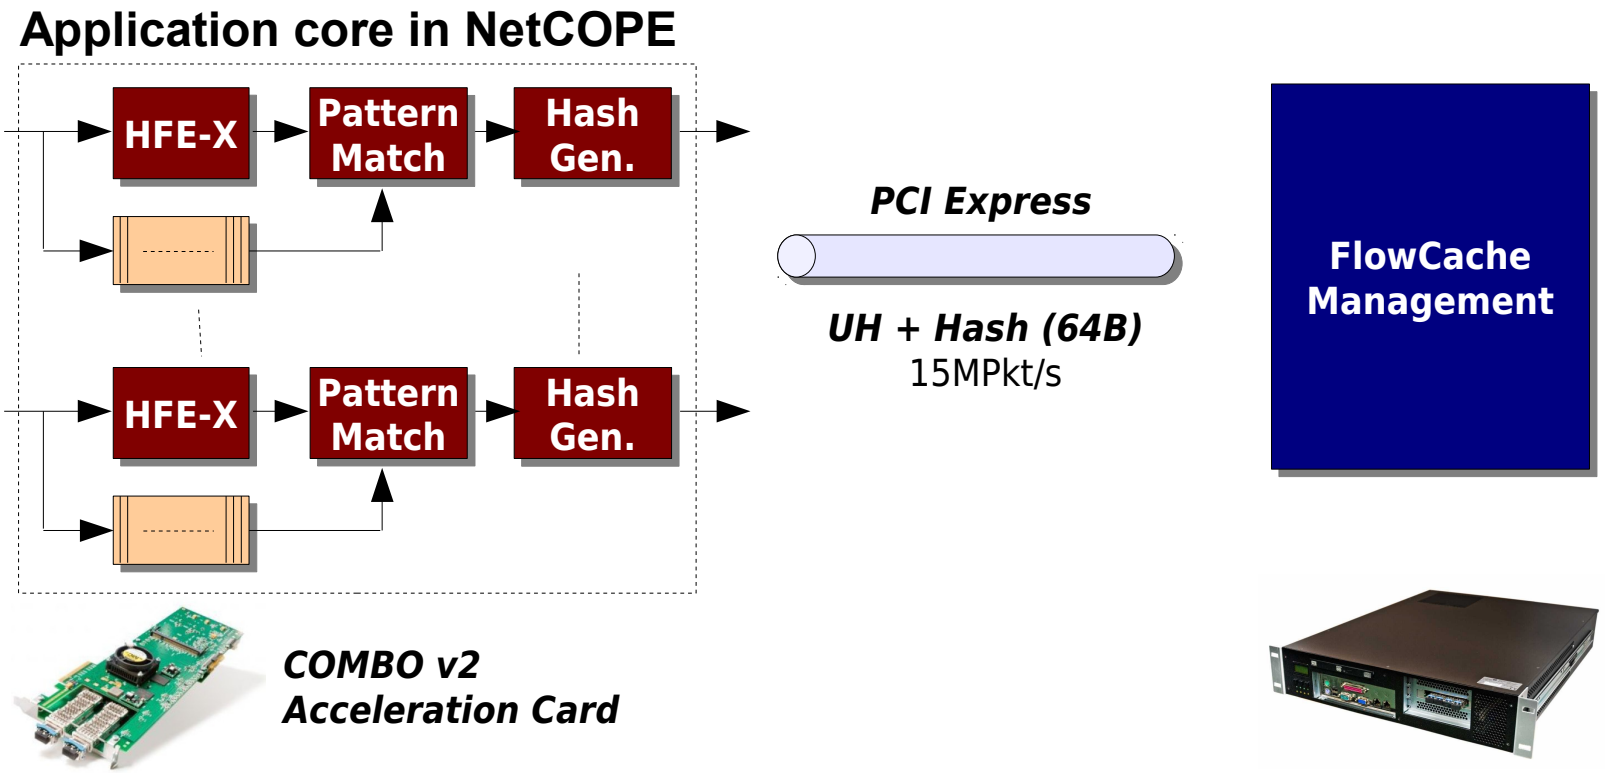

- The probe was designed and implemented in cooperation with CESNET (Czech NREN)

High Speed Probe for Lawful Interception

- Hardware acceleration of packet filtering based on IP addresses, Flows and protocols
- Configuration and management in a host computer
- Designed for wire speed throughput for two 10 Gbps ports

Processing pipeline in NetCOPE platform

COMBO v2

Acceleration card

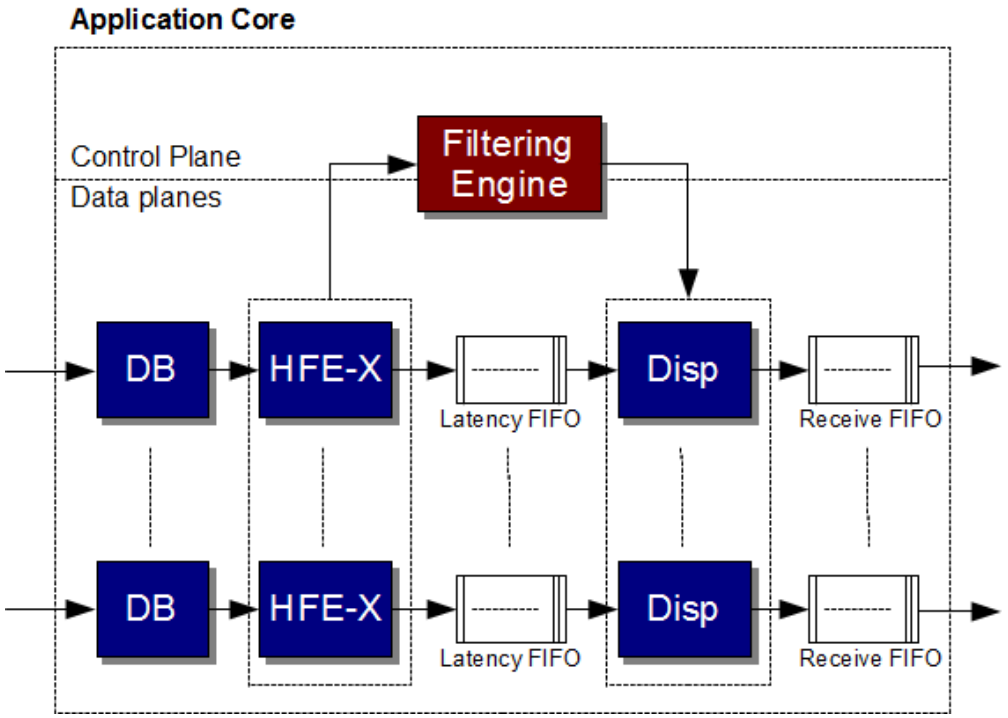

- The probe was designed for Ministry of Interior in CR

Conclusions

- Hardware acceleration in computer networks is necessary for many applications
  - network security, network monitoring and lawful interception, precise packet generator, etc.
- We focus on time-critical operations and hardware acceleration for 40 and 100 Gbps networks
  - Packet header analysis and header fields extraction (up to 40Gbps)
  - Longest prefix matching with Hash and TreeBitmap is ready for 100Gbps even for IPv6 networks
  - Packet classification with perfect hash crossproduct algorithm can achieve 100Gbps throughput with only two QDR SRAM
  - Pattern matching with NFA Split architecture is only for 10Gbps networks, but reduce required logic to ½
- NetCOPE Framework enables rapid prototyping of new applications
  - High speed DMA throughput enables to split application between hardware and software
  - Prepared building blocks enables to create processing pipeline very fast
- Most of the technology have been transferred to INVEA-TECH company which is Brno University of Technology spin-off
- We cooperate on target applications with CESNET, Ministry of Interior (Czech Police) and INVEA-TECH