High Throughput and Programmable Online Traffic Classifier on FPGA

Da Tong, Lu Sun, Kiran Kumar Matam, Viktor K. Prasanna

Ming Hsieh Department of Electrical Engineering

University of Southern California

Evolving Internet

- High-speed packet forwarding
  - Growing network/application demands
  - 10/40 Gbps → 100/400 Gbps
  - E.g. NTT Communications 600 Gbps link Japan-USA

- Network management
  - Flow prioritization
  - Traffic shaping
  - Traffic policing

- Network security
  - Filter/block network traffic/attacks
  - Application level security
  - Firewalls, access control lists, etc.

Based on accurate traffic classification

Traffic Classification at Flow Level (1)

- Determine the application protocol of a traffic flow by inspecting its content.
  - Traffic flow:
    A series of packets sharing the same 5 tuple information within a time window
  - 5 tuple information:
    {Source IP, Destination IP, IP Protocol, Source Port, Destination Port}
    For example the 5 tuple information of an HTTP packet:

262.154.23.2	115.114.35.63	TCP	11689	80

  - What content to inspect?
    - Header information
    - Packet payloads
    - Connection patterns
    - …

Traffic Classification at Flow Level (2)

- Existing techniques
  - Payload based
    - Inspect application layer payload

    ✗ Encrypted payload
  - Port number based
    - Inspect source and destination port numbers

    ✗ Dynamic port assignment
  - Heuristic based
    - Inspect connection patterns

    ✗ Low accuracy & large memory requirement
  - Machine learning based
    - Inspect statistical properties of flows

    ○ Accurate & robust

Machine Learning based Traffic Classification

- Uses statistical properties of the application protocol
  - Statistical properties are referred to as “flow features”
  - Max./Min./Avg. packet size/packet inter-arrival time
  - Port numbers, …
- Off-line training + On-line classification
- Highly accurate if
  - The training data is accurate
  - Proper features are used
- C4.5 Decision tree
  - Well know machine learning technique
  - Highly accurate with various target applications, test traces, and experimental setups in the previous works

Problem Definition

- Design a C4.5 decision tree based traffic classifier on FPGA
  - Assumption: a preceding system will compute the feature vector
  - Input: feature vectors of the flow
  - Output: application protocol of the input flow

- Goals
  - High accuracy:
    - >90% true positive rate

    ← Feature selection using Internet traces from major ISP
  - High throughput:
    - >400 Gbps

    ← Deep pipelining & multi-threaded design
  - Programmability:
    - Support various C4.5 models

    ← Programmable memory structure

Main Contributions

- Identified appropriate features for high accuracy traffic classification
  - Can classify traffic traces consisting of 8 major application protocols
  - Empirically optimized feature set
  - 97.92% overall true positive rate

- Designed programmable architecture
  - Programmable memory structure
  - Extensible to handle updates for decision tree model at run-time

- Designed high throughput architectures on state of the art FPGA
  - 550 Gbps for Dist. RAM based pipelined design
  - 449 Gbps for Block RAM based multi-threaded design

Feature Selection (1)

- Criterion for candidate feature
  - High discriminative power
  - Low computational cost
  - Early classification

- Candidate features
  - IP protocol
  - Src. port number
  - Dst. port number

  } Classic features

  - Sizes of the first N packets
  - Avg./Max./Min. packet size of the first N packets
  - Var. of packet size of the first N packets
  - N = 1,2,3,…8

  } Statistical features

Feature Selection (2)

- Methodology
  - Combine candidate features to construct feature sets
  - Construct C4.5 decision trees using different feature sets
  - Compare their accuracy over all the applications

- Application Protocol
  - HTTP
  - MSN
  - P2PTV
  - QQ_IM
  - Skype
  - Skype_IM
  - Thunder
  - Yahoo_IM

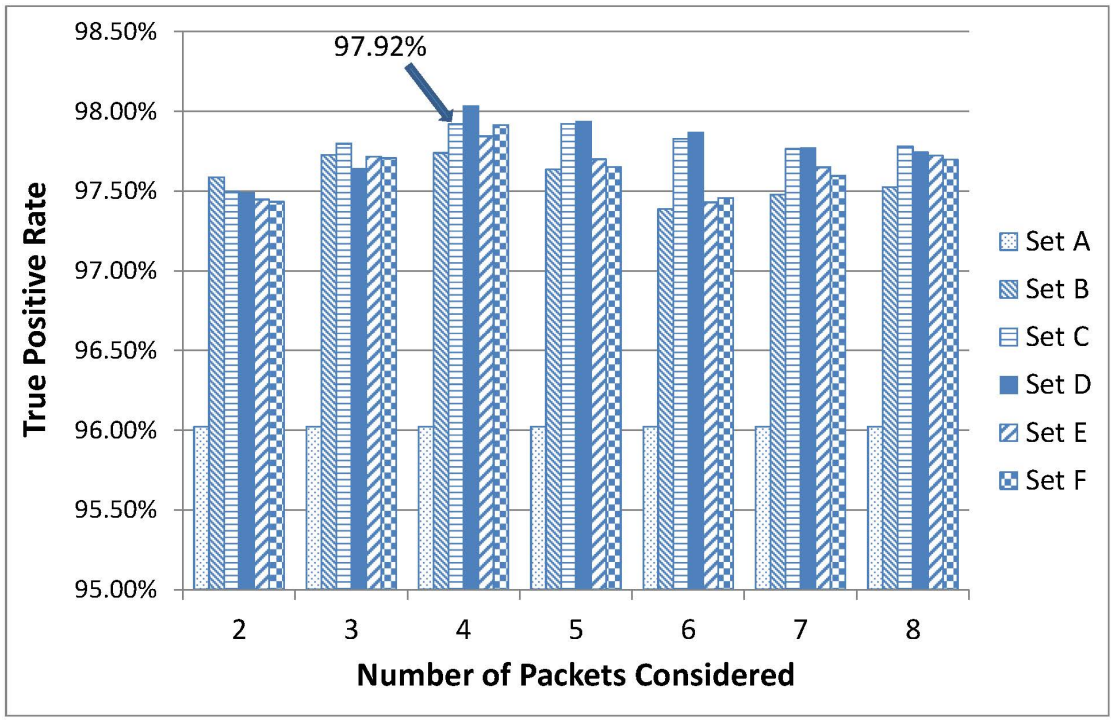

Feature Selection (3)

- The empirically optimized feature set
  - For a mixed traffic trace consisting of 8 application protocols
  - Classic features: IP protocol, src. port number, and dst. port number
  - Statistical features: avg., max., and min. packet size of the first 4 data packets

- Both classic and statistical features are necessary
  - Classic features distinguish classic applications
    - Loss of over 10% accuracy if not included
  - Statistical features distinguish P2P applications
    - Loss of over 1% - 8% accuracy in classifying P2P applications if not included

- Variance is excluded due to high computational cost
  - Loss of only ~0.1% overall accuracy
  - High logic and storage requirement due to square operation

Programmable Architecture

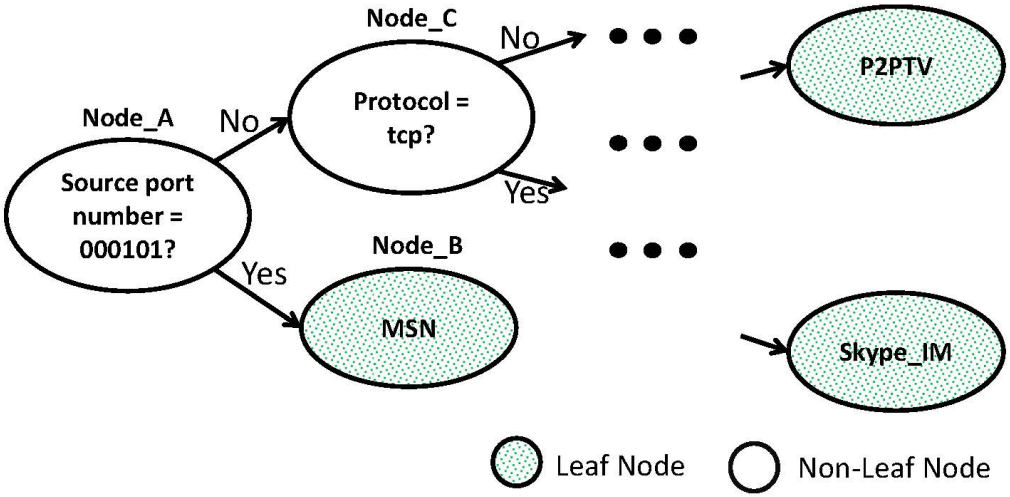

- Node data → memory structure
  - Data can be reprogrammed to support various tree models
- Operation → logic
  - No compilation needed when model changes
- Able to support various C4.5 models

High Throughput Architecture

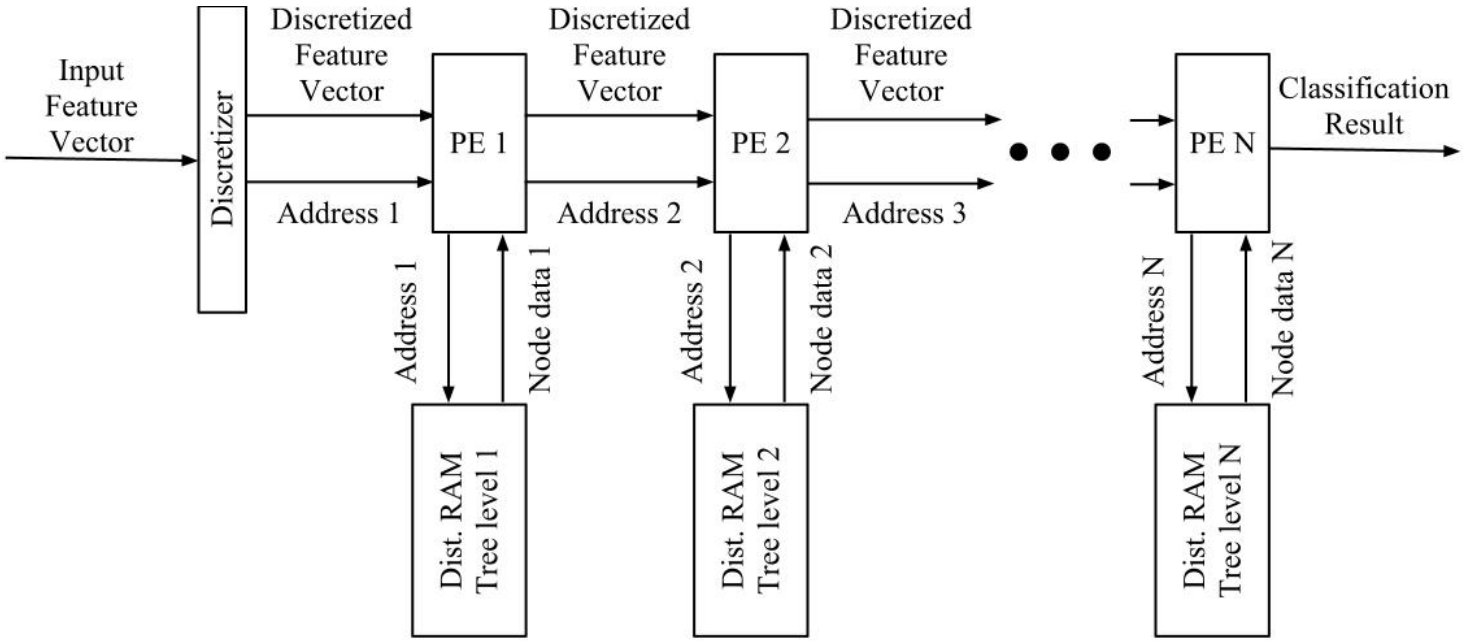

- Localized distributed RAM
  - Distributed RAM block is close to its processing element
  - Low routing requirement → High clock rate
- Deep pipelining
  - Classic approach to achieve high throughput

Multi-threaded Architecture

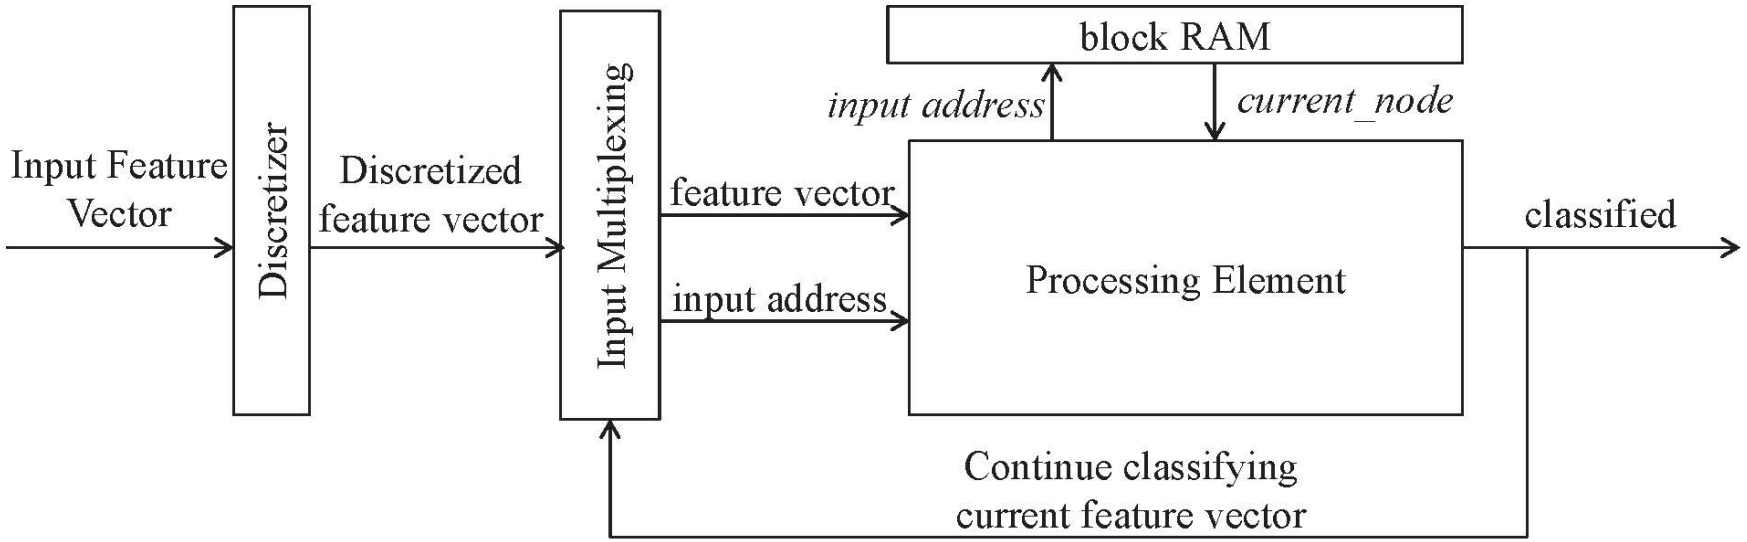

- Localized PE
  - PE is close to BRAM→ Low routing requirement → High Clock Rate
  - Highly scalable
- Multi-threaded parallelism
  - Could be a good approach if memory requirement is small

Implementation

- Xilinx Virtex 6 VLX760
- Dual-port RAM on FPGA
  - Each RAM block serves two pipeline stages/threads
- Deep pipeline design & multi-threaded design
  - Explore both type of parallelism to achieve high throughput

Throughput

- Classifier model
  - Most accurate model using the empirically optimized feature set
  - 43 levels
  - No more than 6 nodes per level

- High throughput design
  - Clock rate: 215 MHz
  - 1 flow/cycle, 4 packets/flow, 40 bytes/packet
  - Throughput: 550 Gbps

- Multi-threaded design
  - Clock rate: 308 MHz
  - 43 cycles/flow
  - Throughput: 6 Gbps/thread
  - Highly scalable when the memory requirement is small
    - For a tree of size 1024 nodes, 449 Gbps by using 160 threads

Scalability of Multi-threaded Architecture

- Tree models of various sizes

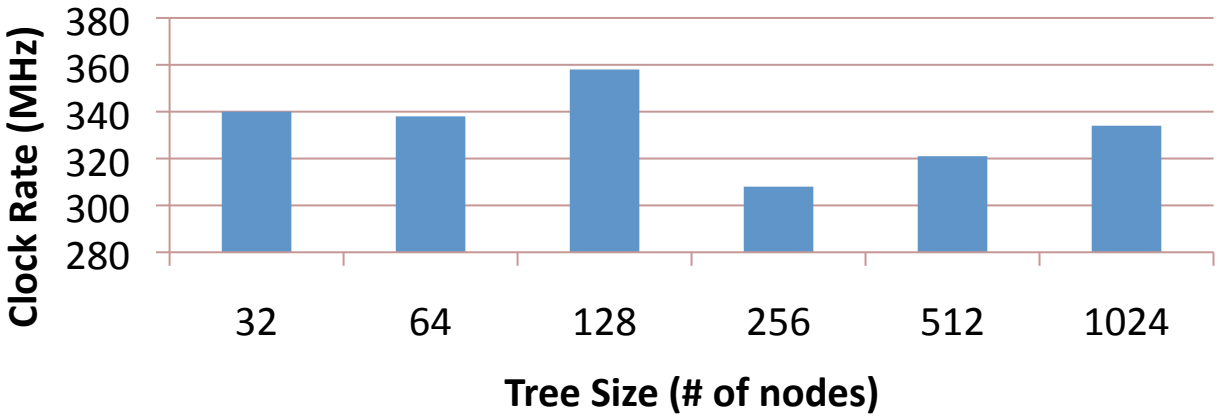

- Various number of threads

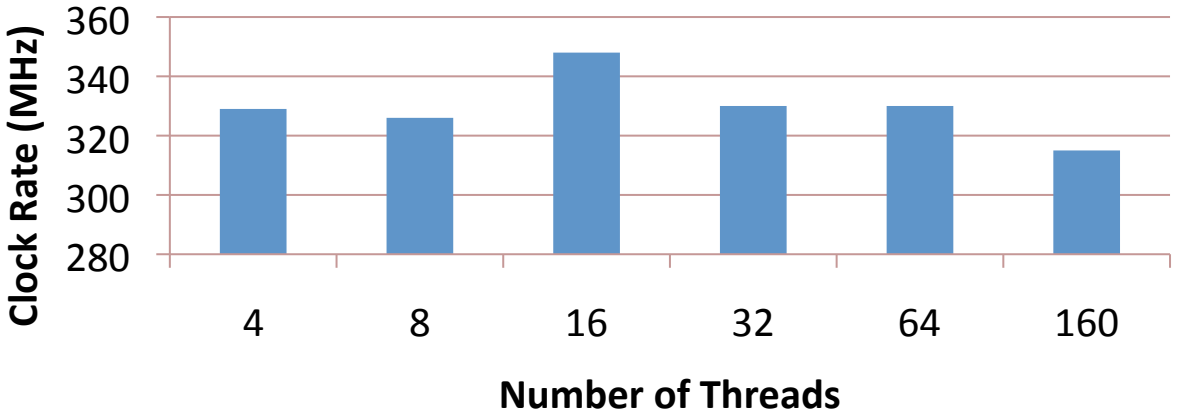

Summary and Future Work

- What we have achieved
  - Identified appropriate features for high accuracy traffic classification
  - Designed programmable architecture to support various C4.5 decision tree models
  - Designed the first 400 Gbps single chip traffic classifier
    - Both deep pipelining and multi-threaded parallelism have been explored

- Future Work
  - Dynamic updating of the C4.5 model in both our architectures
  - Explore the potential of the multi-threaded parallelism in high throughput network processing applications

Thank you!

Questions?