IP lookup on FPGA

Swapnil Haria
Faculty Advisor : Viktor Prasanna
Graduate Students : Thilan Ganegedara, Yun Qu, Edward Yang

USC Viterbi School of Engineering

Routers and IP Lookup

Routers

Forward incoming packets toward their final destinations

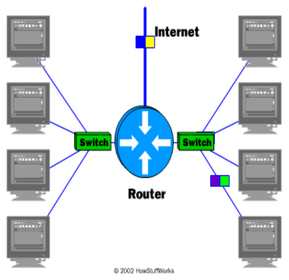

Routing Table

Mapping between IP address -> Forwarding information

010001010.... Destination IP

Prefix	Next Hop
0*	N3
000*	N1
001*	N2
010*	N1
0100*	N4
0101*	N3

The IP Lookup Problem

To find the longest prefix in the table that matches the packet destination address

Router Virtualization

- **Concept** Single hardware router serves packets from multiple networks.
- **Challenge** Ensuring the same throughput and Quality of Service (QoS) as non-virtualized (NV) routers.
- **Benefits** Eases administrative issues of maintaining multiple networking hardware.
  More Throughput per unit power consumption!

FPGA is the ideal candidate!
+Parallelism
+Improved Power Efficiency
-Hardware Resources

Two types :- Virtualized Merged (VM), Virtualized Separate (VS)

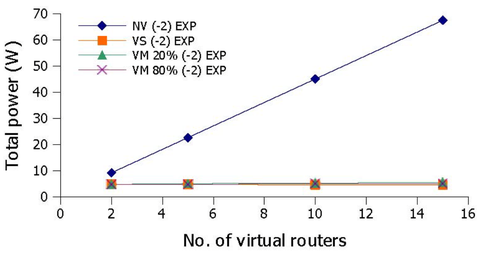

Power Performance of Routers

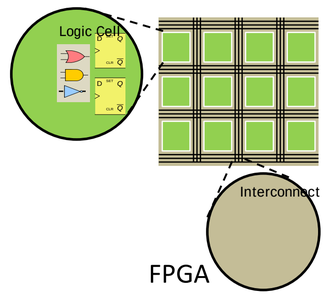

Background : Trie-Based IP Lookup

Uni-bit Trie : Common data structure for storing Lookup Table

Properties

- prefix is represented by a Node
- value of the prefix corresponds to the path from the root of the tree to the prefix node
- Leaf nodes store next hop information in a vector corresponding to each virtualized network

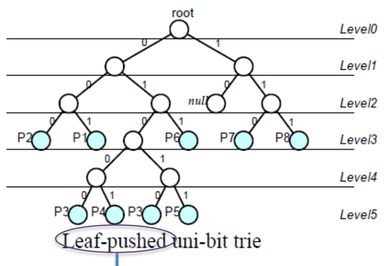

All prefixes at leaf nodes only

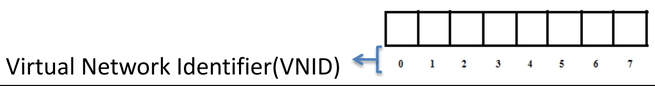

Virtualization Models

Virtualized Separate Model

- Separate routing tables and look-up engines maintained
- Spatial sharing of hardware only.

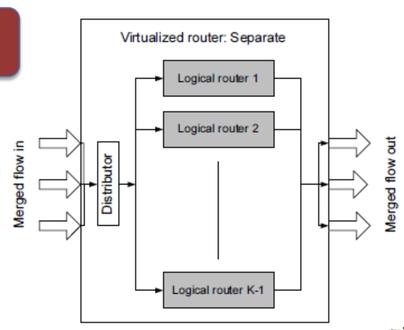

Virtualized Merged Model

- Routing tables merged into a single table.
- Time sharing of the hardware.

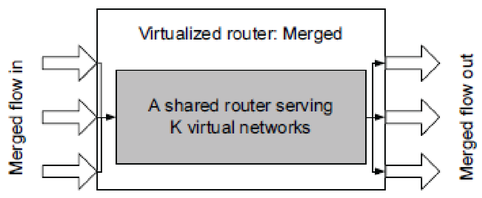

GOAL

To develop a Power-Efficient and Scalable Virtual Router Architecture on FPGA

Complete Architecture

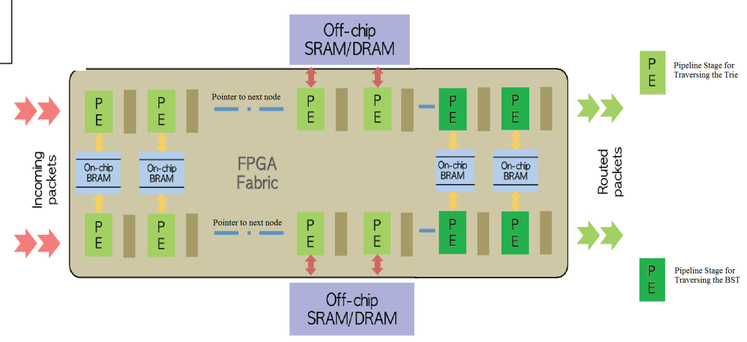

Optimizations to the VM model

We propose three optimizations to the VM model -

1. Replacing Leaf Node Vectors with BSTs
2. Separate storage of small prefixes
3. Selective enabling of Memory

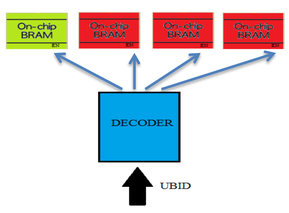

Reduces dynamic power by turning off unused memory (3)

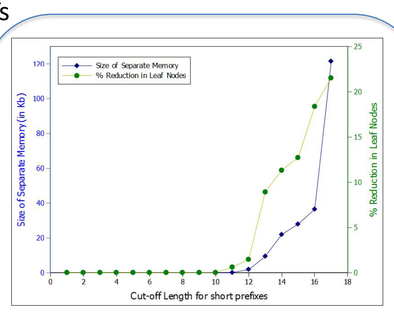

Exploits the fact that shorter prefixes are few, but lead to more leaf nodes on leaf pushing (2)

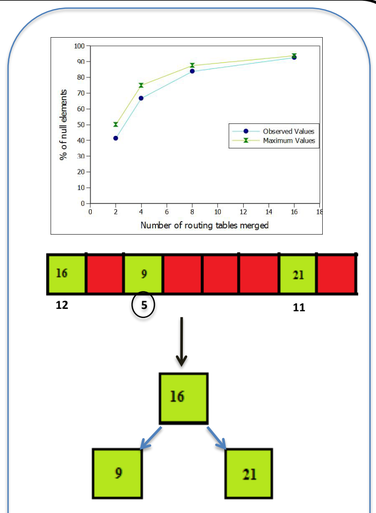

Exploits the low overlap between merged routing tables (1)

Experimental Results

Nodes	Memory		BRAM Power	
	Optimized	Unoptimized	Optimized	Unoptimized
Non-leaf	35.65Mb	35.65Mb	0.073 W	2.273 W
Leaf	13.5 Mb	30Mb	0.006 W	1.967 W
Overall	49.15 Mb	65.65 Mb	0.079 W	4.24 W

Comparison of Memory and Power Consumption between optimized VM model and the unoptimized (conventional)

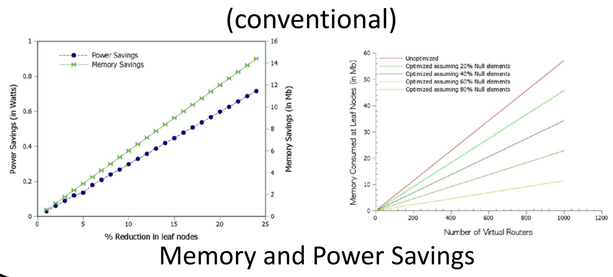

Memory and Power Savings

Major Contributions

- Modified Virtualized Merged Model
  - ✓Reduced Power Consumption
  - ✓Reduced Memory Consumption
  - ✓Scalable
- Analysis of real world routing tables to understand the memory distribution on merging the routing tables.

Scalable Unified Lookup Architecture for IPv4 and IPv6 on FPGA

Salient Features of the proposed Architecture-

- **Unified** process for both IPv4 and IPv6 lookup
- **Scalable**, with respect to routing table size
- **Flexible** design time trade-offs possible between throughput, routing table size, and power usage
- **Robust**, no dependence on prefix distribution and results reported for worst case memory and power consumption

Further Contributions

Mathematical Model developed to illustrate the relationship between performance metrics, resource constraints and design parameters.

Complete Architecture

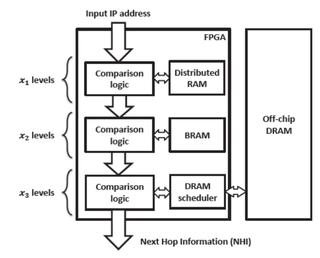

Comparison with other approaches

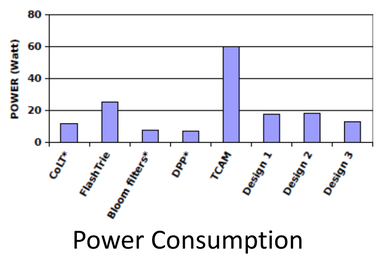

Power Consumption

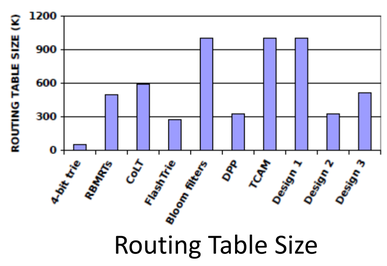

Routing Table Size

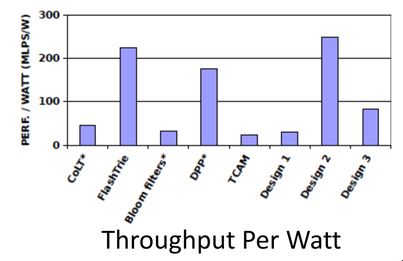

Throughput Per Watt

Results

Three implementations shown
- each focuses on a particular performance metric.

Design Type	Table Size	Clock Rate (MHz)	Throughput (MLPS)	Power (W)
large-table	~1M	275.482	550.964	17.771
high-throughput	~330k	254.518	1527.108	18.587
power-saving	~512k	282.326	564.652	11.198

Performance-optimized mapping of multiple packet lookup schemes on FPGA

Context

Ethernet/IP based packet forwarding has transformed into a complex set of lookup schemes.

Problem

Manual conversion of lookup schemes into FPGA-based design
- time-consuming activity ,
- does not ensure the optimal allocation of FPGA resources ,
- high packet latency.

Solution

We propose a pure Integer Linear Programming (ILP) model for optimized hardware resource planning in network applications.

Results

Lookup Schemes	ILP solution			Manual		
	On-chip mem.	I/O pins	Latency	On-chip mem.	I/O pins	Latency
Ethernet Bridging I	15.97	148	11	15.97	148	11
Ethernet Bridging II	16.6656	360	13	17.4	508	14
Mac-in-Mac	19.4304	506	13	19.5	792	17
Q-in-Q	19.4304	506	13	19.5	960	17

RESOURCE UTILIZATION COMPARISON FOR MULTIPLE LFG MAPPING

Major Contributions

- Pure ILP formulation for the mapping problem
- High-Level model for FPGA based implementations
- MILP formulation for power-efficient mapping

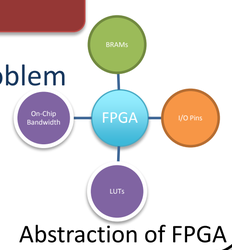

Abstraction of FPGA

Outcome

Power-Efficient and Scalable Virtual Router Architecture on FPGA
-S. Haria, T. Ganegedara and V. Prasanna
-Scheduled for submission to ReConFig'12

Performance-optimized mapping of multiple packet lookup schemes on FPGA
-S. Haria, V. Prasanna
-Scheduled for submission to IEEE INFOCOM 2013

Scalable Lookup Architecture for IPv4 and IPv6 on FPGA
-Y. Qu, S. Haria, Y. Yang, V. Prasanna
-Scheduled for submission to IEEE INFOCOM 2013