Review of:

Packet Classification Using Multidimensional Cutting

- Paper by:
  - Sumeet Singh, Florin Baboescu, George Varghese (UC San Diego)
  - Jia Wang (AT&T Labs)
- Published in:
  - SIGCOMM 2003 (Karlsruhe, Germany)
  - August 25-29, 2003
- Review by:
  - Sarang Dharmapurikar

Motivation

- Importance
  - QoS and Security
- Performance of existing schemes
  - RFC : fast but excessive storage
  - EGT-PC : speed still slow for practical purposes
  - HiCuts : same..
  - TCAM : high power consumption, high cost and less dense
    - Vendors like Cypress, Fast-Chip, EZchip, Integrated Silicon looking for different alternatives
- New ideas
  - New ideas still possible
  - Authors introduce one such idea : HyperCuts

Packet Classification Problem

$Rule$	$Field_1$	$Field_2$	$Field_3$	$Field_4$	$Field_5$	$ACTION$
R_0	000*	111*	10	*	UDP	act_0
R_1	000*	111*	01	10	UDP	act_0
R_2	000*	10*	*	10	TCP	act_1
R_3	000*	10*	*	01	TCP	act_2
R_4	000*	10*	10	11	TCP	act_1
R_5	0*	111*	10	01	UDP	act_0
R_6	0*	111*	10	10	UDP	act_0
R_7	0*	1*	*	*	TCP	act_2
R_8	*	01*	*	*	TCP	act_2
R_9	*	0*	*	01	UDP	act_0
R_{10}	*	*	*	*	UDP	act_3
R_{11}	*	*	*	*	TCP	act_4

$Rule$	$Field_1$	$Field_2$	$Field_3$	$Field_4$	$Field_5$	$ACTION$
R_0	0 – 1	14 – 15	2	0 – 3	0	act_0
R_1	0 – 1	14 – 15	1	2	0	act_0
R_2	0 – 1	8 – 11	0 – 3	2	1	act_1
R_3	0 – 1	8 – 11	0 – 3	1	1	act_2
R_4	0 – 1	8 – 11	2	3	1	act_1
R_5	0 – 7	14 – 15	2	1	0	act_0
R_6	0 – 7	14 – 15	2	2	0	act_0
R_7	0 – 7	8 – 15	0 – 3	0 – 3	1	act_2
R_8	0 – 15	4 – 7	0 – 3	0 – 3	1	act_2
R_9	0 – 15	0 – 7	0 – 3	1	0	act_0
R_{10}	0 – 15	0 – 15	0 – 3	0 – 3	0	act_3
R_{11}	0 – 15	0 – 15	0 – 3	0 – 3	1	act_4

Graphical Illustration

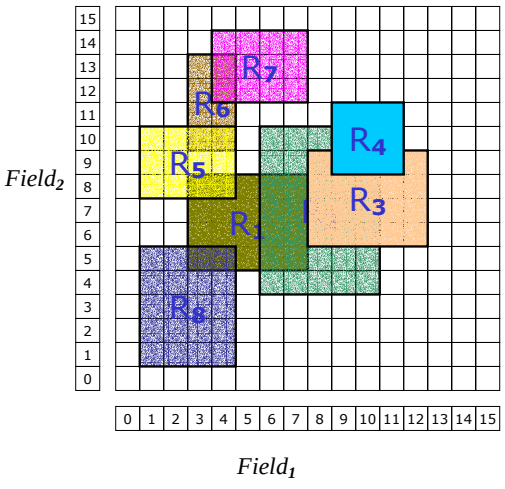

$Field_1$	$Field_2$	$Rule$
3-7	5-8	R_1
6-10	4-10	R_2
8-12	6-9	R_3
9-11	9-11	R_4
1-4	8-10	R_5
3-4	10-13	R_6
4-7	12-14	R_7
1-4	1-5	R_8

Constructing a Decision Tree

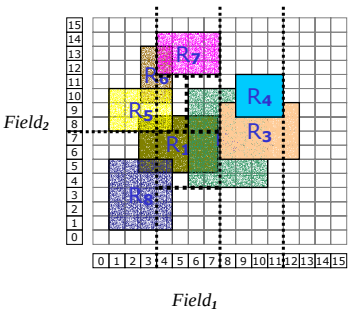

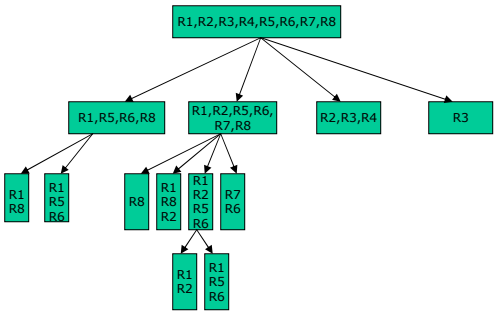

- Goals
  - Tree with minimal depth
  - With least nodes
  - With uniformly distributed rules in the leaves
  - With reasonably small # of rules in the leaves
- Questions
  - How many dimensions to cut at a time?
  - Which dimensions to choose?
  - How many cuts on which dimension?

Hierarchical Cuts (HiCuts)

- How many fields to cut?
  - Only one field on one node
- Which field to cut?
  - Some heuristics
    - field with maximum unique components specified
    - Minimize the maximum # of rules in a child node
- How many cuts to make on it?
  - Some heuristics
    - Total # cuts at each node bounded by *spmf(n) = spfac x n*

HyperCuts Illustration

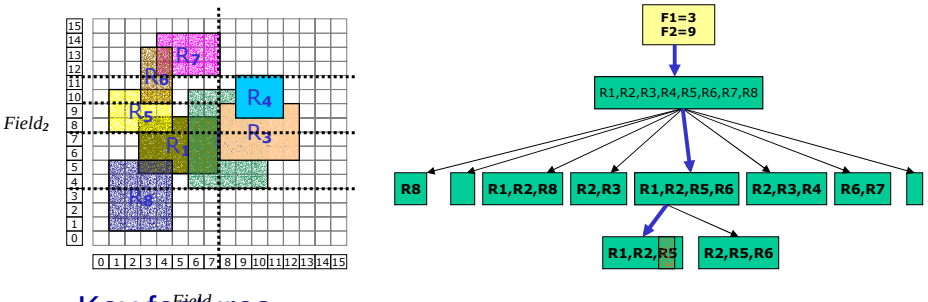

- Key features
  - Tree depth is reduced by cutting multiple fields
  - Navigation through tree is done by array indexing

HyperCuts

- How many fields to cut?
  - One or more
- Which fields to cut?
  - Some heuristics
    - Field with number of unique elements greater than the mean of the number of unique elements across all the dimensions considered
    - Also experimented with ratio of number of unique elements to the size of that region across the dimension chosen
  - Denote the set of fields to cut on a node by $\boldsymbol{D}$
- How many cuts to make on each field?
  - Some heuristics
    - Cuts on field i = $nc(i)$ $i \in D$, decided by some heuristics
    - Total cuts at one node = $\Pi\, nc(i) < spfac \times \sqrt{N}$

Optimizations

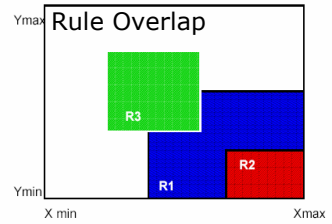

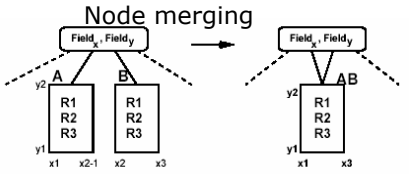

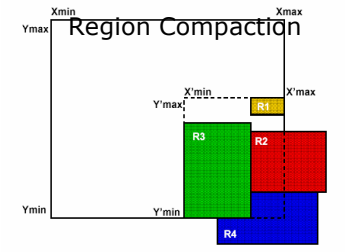

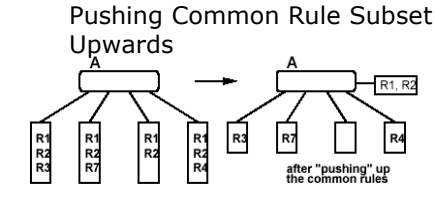

Results on Core Router Databases

Memory space requirement

Database	*No.Rules*	*RFC*	*HiCuts* − 4	*HiCuts* − 1	*ABV*	*EGT* − *PC*	*HypCuts* − 4	*HypCuts* − 1
CR_1	85	55,202	11,608	1,346	1,572	1,168	453	226
CR_2	125	114,080	10,704	1,986	1,606	1,472	589	610
CR_3	351	100,991	64,541	19,001	4,651	2,261	15,395	11,210
CR_4	2799	747,271	117,801	25,543	285,099	30,753	16,631	11,030

Observed worst case memory accesses

Database	*No.Rules*	*RFC*	*HiCuts* − 4	*HiCuts* − 1	*ABV*	*EGT* − *PC*	*HypCuts* − 4	*HypCuts* − 1
CR_1	85	12	25	32	111	32	14	14
CR_2	125	12	25	36	106	54	8	11
CR_3	351	12	35	57	126	47	22	31
CR_4	2,799	12	38	66	196	87	18	25

Results on Edge Router Databases

Memory space requirement

Database	*No. of Rules*	$EGT-PC$	$HiCuts-4$	$HiCuts-1$	$HyperCuts-4$	$HyperCuts-1$
ER_1	4740	284,159	6695	6659	6659	6482
ER_2	2505	149,470	3730	3470	3456	3393
ER_3	995	62,266	1501	1459	1527	1465
ER_4	2458	154,976	3263	3274	3295	3295

Observed worst case memory accesses

Database	*No. of Rules*	$EGT-PC$	$HiCuts-4$	$HiCuts-1$	$HyperCuts-4$	$HyperCuts-1$
ER_1	4740	62	15	18	15	18
ER_2	2505	63	15	18	15	18
ER_3	995	49	15	18	15	15
ER_4	2458	65	15	15	11	15

Results on Firewall Databases

Memory space requirement

Database	*No. of Rules*	$EGT-PC$	$HiCuts-4$	$HiCuts-1$	$HyperCuts-4$	$HyperCuts-1$
FW_1	279	7,477	48,347	16,978	9,574	6,026
FW_2	183	3,642	20,995	2,872	4,311	6,675
FW_3	158	2,962	18,207	6,675	2,164	943
FW_4	266	4,275	14,624	6,375	9,477	6,991

Observed worst case memory accesses

Database	*No. of Rules*	$EGT-PC$	$HiCuts-4$	$HiCuts-1$	$HyperCuts-4$	$HyperCuts-1$
FW_1	279	63	41	74	26	32
FW_2	183	55	74	74	20	26
FW_3	158	56	74	74	17	17
FW_4	266	38	23	50	17	23

Results on Synthetic Databases

Database	*No. of Rules*	*Memory Space*	*Search*
$ER-5K$	4,970	6,745	14
$ER-10K$	9,940	13,482	29
$ER-15K$	14,910	20,188	32
$ER-20K$	19,880	27,098	35
$CR-5K$	5,000	4,508	8
$CR-10K$	10,000	8,495	8
$CR-15K$	15,000	12,822	11
$CR-20K$	20,000	16,857	11
$FW-5K$	5,000	14,733	14
$FW-10K$	10,000	45,158	14
$FW-15K$	15,000	80,486	17
$FW-20K$	20,000	150,551	17