



Chad R. Meiners
Alex X. Liu
Eric Torng

Hardware Based Packet Classification for High Speed Internet Routers

 Springer

Hardware Based Packet Classification for High Speed Internet Routers

Chad R. Meiners • Alex X. Liu • Eric Torng

Hardware Based Packet Classification for High Speed Internet Routers

 Springer

Chad R. Meiners
Michigan State University
Department of Computer Science
and Engineering
Engineering Building 3115
East Lansing, MI 48824-1226
USA
meinersc@cse.msu.edu

Alex X. Liu
Michigan State University
Department of Computer Science
and Engineering
Engineering Building 3115
East Lansing, MI 48824-1226
USA
alexliu@cse.msu.edu

Eric Tornø
Michigan State University
Department of Computer Science
and Engineering
Engineering Building 3115
East Lansing, MI 48824-1226
USA
tornø@msu.edu

This material is based upon work supported by the National Science Foundation under Grant No. 0916044. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

ISBN 978-1-4419-6699-5 e-ISBN 978-1-4419-6700-8
DOI 10.1007/978-1-4419-6700-8
Springer New York Dordrecht Heidelberg London

Library of Congress Control Number: 2010928719

© Springer Science+Business Media, LLC 2010

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

This book is dedicated to chance encounters at the correct time.
– Chad Meiners

*Dedicated with love and respect
to my parents Shuxiang Wang and Yuhai Liu (God rest his soul),
to Huibo Heidi Ma
to my twin sons Max Boyang and Louis Boyang,
to whom I owe all that I am and all that I have accomplished.*
– Alex X. Liu

Dedicated to Pat and Annika.
– Eric Torng

Contents

1	Introduction	1
1.1	Motivation	2
1.2	Contribution	3
2	Background	5
2.1	Firewall decision diagrams	6
2.2	One-Dimensional Classifier Minimization	9
3	Related Work	15
3.1	TCAM Based Classifiers	15
3.1.1	Classifier Minimization:	15
3.1.2	Range Encoding:	18
3.1.3	Circuit Modification:	20
3.2	Software Based Techniques	20
3.2.1	Parallel decomposition	21
3.2.2	Decision Trees	23

Part I Equivalent Transformation Techniques

4	TCAM Razor	31
4.1	Multi-dimensional TCAM Minimization: The Basics	31
4.1.1	Conversion to Firewall Decision Diagrams	31
4.1.2	Multi-dimensional TCAM Minimization	32
4.1.3	Removing Redundant Rules	34
4.1.4	The Algorithm	34
4.1.5	TCAM Update	35
4.2	Multi-dimensional TCAM Minimization: The Optimization Techniques	35
4.2.1	Lazy Copying in FDD Construction	35
4.2.2	Hashing in FDD Reduction	36

5	Bit Weaving	41
5.1	Bit Swapping	43
5.1.1	Prefix Bit Swapping Algorithm	43
5.1.2	Minimal Cross-Free Classifier Partitioning Algorithm	45
5.1.3	Partial List Minimization Algorithm	47
5.2	Bit Merging	47
5.2.1	Definitions	48
5.2.2	Bit Merging Algorithm (BMA)	48
5.3	Discussion	51
5.3.1	Redundancy Removal	51
5.3.2	Incremental Classifier Updates	51
5.3.3	Composability of Bit Weaving	52
5.3.4	Prefix Shadowing	53
5.4	Complexity Analysis of Bit Weaving	54
6	All-Match Redundancy Removal	57
6.1	All-Match Based Redundancy Theorem	57
6.1.1	All-Match FDDs	57
6.1.2	The All-Match Based Redundancy Theorem	59
6.2	All-Match Based Redundancy Removal	60
6.2.1	The All-Match FDD Construction Algorithm	60
6.2.2	The All-Match Based Redundancy Removal Algorithm	60
6.2.3	Proof of Complete Redundancy Removal	63
6.3	Optimization Techniques	64
6.3.1	Decision Chaining	64
6.3.2	Isomorphic Terminal Nodes Elimination	66
6.4	Classifier Incremental Updates	67
6.4.1	Rule Insertion	68
6.4.2	Rule Deletion	68
6.4.3	Rule Modification	69
6.5	Redundancy Analysis	69

Part II New Architectural Approaches

7	Sequential Decomposition	75
7.1	Multi-Lookup Approach	77
7.1.1	Constructing Multi-lookup Table	77
7.1.2	Packet Processing	80
7.1.3	Analysis	81
7.2	Pipelined-lookup Approach	82
7.2.1	Pipelined-Table Construction	82
7.2.2	Packet Processing	84
7.2.3	Analysis	84
7.3	Packing Approach	85
7.3.1	Strict Partitioning	86

- 7.3.2 Shadow Packing 89
- 7.3.3 Strict Partitioning vs. Shadow Packing 93
- 7.4 Table Consolidation 93
 - 7.4.1 Table Consolidation Algorithm 93
 - 7.4.2 Hierarchical Table Consolidation 94
 - 7.4.3 TCAM/SRAM Space Tradeoff via Bounded Consolidation . 96
- 7.5 One-Dimensional Table Consolidation 96
 - 7.5.1 Table Consolidation with Sequential Decomposition 96
 - 7.5.2 Coping with More Fields than TCAM chips 97
- 7.6 Implementation Issues 97
 - 7.6.1 TCAM Update 97
 - 7.6.2 Non-ordered FDDs 98
 - 7.6.3 Lookup Short Circuiting 98
 - 7.6.4 Best Variable Ordering 99
- 8 Topological Transformations 101**
 - 8.1 Topological Transformation 103
 - 8.1.1 Architectures 103
 - 8.1.2 Measuring TCAM space 105
 - 8.1.3 TCAM Update 106
 - 8.2 Domain Compression 106
 - 8.2.1 Step 1: Compute Equivalence Classes 107
 - 8.2.2 Step 2: Construct Transformers 107
 - 8.2.3 Step 3: Construct Transformed Classifier 109
 - 8.3 Prefix Alignment 111
 - 8.3.1 Prefix Alignment Overview 111
 - 8.3.2 One-dimensional Prefix Alignment 113
 - 8.3.3 Multi-Dimensional Prefix Alignment 116
 - 8.3.4 Composing with Domain Compression 117
- References 119**
- Index 123**

List of Figures

2.1	Illustration of FDD construction	8
2.2	An example one-dimensional TCAM minimization problem	12
2.3	Illustration of dynamic program	12
4.1	A firewall decision diagram	32
4.2	“Virtual” one-dimensional packet classifier	33
4.3	Lazing copying of subgraphs	36
4.4	Example of lazing copying	37
4.5a	Before FDD reduction	37
4.5b	After FDD reduction	37
5.1	Example of the bit weaving approach	42
5.2	Example of bit-swapping	45
5.3	Applying bit weaving algorithm to an example classifier	45
5.4	Example of Bit Merging Algorithm Execution	50
6.1	A simple packet classifier	58
6.2	An all-match FDD for the packet classifier in Fig 6.1	58
6.3	Constructing an all-match FDD	61
6.4	The containment list and the residency List for the all-match FDD ..	62
6.5	The resulting packet classifier after removing redundant rules	63
6.6	Constructing an all-match tree with decision chaining	66
6.7	All-Match Tree with decision tree chaining constructed in rule order and reverse order.	67
7.1	Reducing information redundancy	76
7.2	FDD generation and reduction	78
7.3	The Multi-lookup scheme	79
7.4	Table generation and table merge in the pipelined-lookup approach	83
7.5	Example of a pipelined-lookup	84

7.6	Strict partitioning of a TCAM chip	87
7.7	Decisions in SRAM	87
7.8	Reassigning the table IDs	88
7.9	Shadow packed tables & shadow packing tree	90
7.10	The shadow packing process.....	92
7.11	Consolidation of 2-dimensional tables	94
7.12	Table consolidation for pipelined-lookup sequential decomposition ..	95
8.1	Example of topological transformations	102
8.2	Multi-lookup architecture	104
8.3	Parallel pipelined-lookup architecture	104
8.4	Chained pipelined-lookup architecture	105
8.5	Step 1 of domain compression	108
8.6	Step 2 of domain compression	109
8.7	Step 3 of domain compression	109
8.8	Example of 1-D prefix alignment	112

List of Tables

1.1	An example packet classifier	1
1.2	TCAM Razor output for the example packet classifier in Table 1.1	1
4.1	A minimum packet classifier corresponding to v_2 in Fig. 4.1	32
4.2	A minimum packet classifier corresponding to v_3 in Fig. 4.1	33
4.3	A minimum packet classifier corresponding to v_1 in Fig. 4.1	34
4.4	Packet classifier generated from the FDD in Figure 4.1	34
4.5	Rules from the FDDs in Figure 4.5a and Figure 4.5b	38
5.1	The second partition benefits from the prefix shadow of the first partition	53
5.2	The output of bit merging determines the input.	54
5.3	The set of all ternary covers generated per round	55

Chapter 1

Introduction

Packet classification, which is widely used on the Internet, is the core mechanism that enables routers to perform many networking services such as firewall packet filtering, virtual private networks (VPNs), network address translation (NAT), quality of service (QoS), load balancing, traffic accounting and monitoring, differentiated services (Diffserv), etc. As more services are deployed on the Internet, packet classification grows in demand and importance.

The function of a packet classification system is to map each packet to a decision (i.e., action) according to a sequence (i.e., ordered list) of rules, which is called a packet classifier. Each rule in a packet classifier has a predicate over some packet header fields and a decision to be performed upon the packets that match the predicate. To resolve possible conflicts among rules in a classifier, the decision for each packet is the decision of the first (i.e., highest priority) rule that the packet matches. Table 1.1 shows an example packet classifier of two rules. The format of these rules is based upon the format used in Access Control Lists on Cisco routers.

Rule	Source IP	Destination IP	Source Port	Destination Port	Protocol	Action
r_1	1.2.3.0/24	192.168.0.1	[1,65534]	[1,65534]	TCP	accept
r_2	*	*	*	*	*	discard

Table 1.1 An example packet classifier

Rule	Source IP	Destination IP	Source Port	Destination Port	Protocol	Action
r_1	1.2.3.0/24	192.168.0.1	0	*	*	discard
r_2	1.2.3.0/24	192.168.0.1	65535	*	*	discard
r_3	1.2.3.0/24	192.168.0.1	*	0	*	discard
r_4	1.2.3.0/24	192.168.0.1	*	65535	*	discard
r_5	1.2.3.0/24	192.168.0.1	[0,65535]	[0,65535]	TCP	accept
r_6	*	*	*	*	*	discard

Table 1.2 TCAM Razor output for the example packet classifier in Table 1.1