An Analysis of Rule-set Databases in Packet Classification

Mahmood Ahmadi, S. Arash Ostadzadeh, and Stephan Wong
Computer Engineering Laboratory
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
{mahmadi, arash, stephan}@ce.et.tudelft.nl

Abstract— **Packet classification has proved to be an important challenge in network processing. It requires comparing each packet against a database of rules and forwarding the packet according to the highest priority matching rule. Packet classification can be seen as the categorization of incoming packets based on their headers according to specific criteria that examine specific fields within a packet header. The criteria are comprised of a set of rules that specify the content of specific packet header fields to result in a match. A packet classifier can be implemented in either software or hardware. An important entity in packet classification algorithms encompasses the rule-set database which includes a set of rules. Each rule is comprised of different fields such as source-IP, destination-IP, source-port, destination-port and protocol fields. A comprehensive and rigorous analysis of these fields appears to be the first systematic step to tackle the problem of packet classification and probably some elaborate optimizations for the existing algorithms. In this paper, we carefully analyze different rule-set databases and present the processed data with illustrative diagrams regarding the distributions of source/destination IPs, source/destination ports and protocol fields in the rule-sets. Based on the extracted information from these databases, we also investigate the outline of hash-based packet classification algorithms and present some trends to optimize this category of packet classification algorithms.**

Keywords: **Packet classification, Tuple space, Hashing, Bloom filter**

I. Introduction

Traditionally, routers forward a packet based on the destination address in it. The support of many different services such as Quality of Service (QoS), Virtual Private Network (VPN), policy-based routing, traffic shaping, firewalls, and network security, increases the importance of packet classification. In order to provide these services, the router must categorize the incoming packets according to different criteria. These criteria are formed based on one or more fields in the packet header. Packet header fields generally include destination and source IP addresses, the protocol type, and the destination and source port numbers as depicted in Figure 1.

Packet classification can be seen as the categorization of incoming packets according to specific criteria that examine specific fields within packet headers. The criteria are comprised of a set of rules that specify the content of specific packet header fields to result in a match [1][2][3].

In this paper, we direct a thorough analysis of different rule-set databases and present the processed data with illustrative diagrams regarding the distributions of the source and destination IPs, the source and destination ports, and protocol fields. The following facts are derived from this study:

1. The valid bits in the source/destinations IP-addresses in the rule-set databases are mostly distributed between bits 0-4 of the first octet and bits 16-32 of the third and forth octets.
2. Specific source port numbers are recognized more than specific destination port numbers in the rule-set databases.
3. Source and destination port extend in the rule-set databases are mostly of large sizes.
4. The number of rules with just a single destination-port is more than their counterpart source-ports in the rule-set databases.

Based on the presented analysis of these rule-set database, we also investigate the hash-based packet classification algorithms and present some trends to optimize this category of packet classification algorithms.

The rest of this paper is organized as follows. Section II describes a summary of related work in this area. Section III gives an account of the rule-set databases and the tuple space algorithms for packet classification. Section IV presents the results' analysis and finally Section V draws the overall conclusions.

II. Related work

In this section, we take a brief look at the previous works regarding the rule-set analysis. In [4], Classbench is introduced as suite tools for benchmarking packet classification algorithms and devices. Classbench includes a rule-set generator which produces synthetic rule-sets that accurately model the characteristics of real rule-sets. The tool suite also includes a trace generator that produces a sequence of packet headers to exercise the synthetic rule-sets. Along with specifying the relative size of the trace, it provides a simple mechanism for controlling locality of reference in the trace. In [5], 793 packet classifiers from 101 different ISPs and enterprise networks and a total of 41,505 rules are collected. Different characteristics are discovered as follows.

1. The classifiers do not contain a large number of rules.

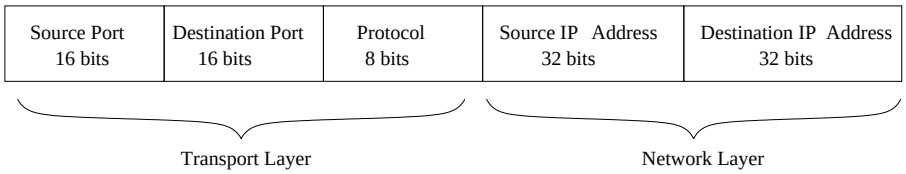

Fig. 1. Fields that are generally used in packet classification algorithms.

2. The syntax allows a maximum of 8 fields to be specified: source/destination network-layer address (32 bits), source/destination transport-layer port numbers (16 bits for TCP and UDP), Type-Of-Service (TOS) field (8 bits), protocol field (8 bits), and transport-Layer protocol flags (8 bits) with a total of 120 bits.
3. The transport-layer protocol field is restricted to a small set of values in all the packet classifiers examined, it contained only TCP, UDP, ICMP, IGMP, (E)IGRP, GRE and IPINIP or the wildcard '*', i.e. the set of all transport protocols.
4. The transport-layer fields have a wide variety of specifications. Many (10.2%) of them are range specifications (e.g. of the type gt 1023, i.e. greater than 1023, or in range 20-24). In particular, the specification 'gt 1023' occurs in about 9% of the rules.
Finally, based on the analysis, a simple multi-stage classification algorithm, called RFC (recursive flow classification) is proposed.

III. Rule-set Database and Tuple Space Packet Classification

In this section, we present the structure of the rule-set databases and the concept of the tuple space classification algorithm.

A. Rule-set Databases

A packet classifier must compare header fields of every incoming packet against a set of rule-sets in order to identify a flow. The resulting flow identifier is used to apply security policies, application processing, and quality-of-service guarantees to packets belonging to the specified flow. Typical packet classification rule-sets have fewer than a thousand rules and reside in enterprise firewalls or edge routers. As network services continue to migrate into the network core, it is anticipated that rule-sets could swell to tens of thousands of rule-sets or more[5]. The most common type of packet classification examines the packet headers fields comprising the standard IP 5-tuple: IP source and destination addresses, transport protocol number, and source, and destination port numbers. A packet classifier searches for the highest priority rule-set or set of rules matching the packet where each rule-set specifies a prefix on the IP addresses, an exact match or wildcard on the transport protocol number, and ranges on the transport port numbers. We utilize 12 real rule-set databases from Applied Research Laboratory in Washington University in St. Louis [4][6] provided by Internet Service Providers (ISPs), a network equipment vendor, and other researchers working in the field.

The rule-set databases range in size from 68 to 4557 entries and utilize one of the following formats:
- Access Control List (ACL) - standard format for security, VPN, and NAT rule-sets for firewalls, and routers (enterprise, edge, and backbone).
- FireWall (FW) - proprietary format for specifying security rule-sets for firewalls.
- IP Chain (IPC) - decision tree format for security, VPN, and NAT rule-sets for software-based systems.

In rule-set databases, each rule combines 5 tuple defined as "*[*Source IP address, Destination IP address, Source port, Destination port, Protocol]" and the format is "*@[Source IP address prefix in dot-decimal notation]/[Prefix length] [Destination IP address prefix in dot-decimal notation]/[Prefix length] [Low source port] : [High source port] [Low destination port] : [High destination port] [Protocol value in hexadecimal]/[Protocol mask in hexadecimal]*" [6]. As an example, a rule in the rule-set database is defined like: *@204.152.188.80/28 204.152.188.64/28 67 : 67 67 : 67 0x11/0xff*.

The packet header trace format is "*[Source IP address in decimal] [Destination IP address in decimal] [Source port value in decimal] [Destination port value in decimal] [Protocol in decimal]*". As an example, a packet header in packet trace might look like: *3337533518 2390673931 65535 65535 1 9*. The specifications of rule-set databases and packet traces are shown in Table I. Table I includes seven rule-set databases and packet traces. Note that the rule-sets and packet traces are one-to-one correspondent. The rule-sets *fw1*, *acl1*, and *ipc1* are extracted from real rule-sets and the others are generated by Classbench benchmark [6].

B. Tuple Space Classification

A high-level approach for multiple field search employs tuple spaces with a tuple representing information in each field specified by the rules. Srinivasan, et al., [7][8], introduced the tuple space approach and the collection of tuple search algorithms. We provide a simplified example of rule classification on five fields in Table II. Address prefixes cover 32-bit addresses and port ranges cover 16-bit port numbers. For address prefix fields, the number of specified bits is simply the number of non-wildcard bits in the prefix. For

Rule Database	fw1-100	fw1-1k	fw1-5k	fw1-10k	fw1	acl1	ipc1
Number of rules	92	971	4653	9311	266	752	1550
Number of tuples	26	42	52	57	36	44	179
Packet trace	**fw1-100**	**fw1-1k**	**fw1-5k**	**fw1-10k**	**fw1**	**acl1**	**ipc1**
Number of Packets	920	8050	46700	93250	2830	8140	17020

TABLE I
RULE-SET DATABASES AND PACKET TRACE SPECIFICATIONS.

Rules	Destination IP (address mask)	Source IP (address mask)	Port No.	Protocol No.	Tuple space
R1	192.168.190.69 (255.255.255.255)	192.168.80.11 (255.255.255.0)	*	*	[32,32,0,0]
R2	192.168.3.0 (255.255.255.0)	192.168.200.157 (255.255.255.255)	eq www	tcp	[24,32,2,1]
R3	192.168.198.4 (255.255.255.255)	192.168.160.0 (255.255.255.0)	gt 1023	udp	[32,32,1,1]
R4	193.164.0.0 (255.255.0.0)	193.0.0.0 (255.0.0.0)	eq www	udp	[16,8,2,1]
R5	192.168.0.0 (255.255.0.0)	192.0.0.0 (255.0.0.0)	eq www	tcp	[16,8,2,1]
R6	0.0.0.0 (0.0.0.0)	0.0.0.0 (0.0.0.0)	*	*	[0,0,0,0]

TABLE II
SIMPLIFIED EXAMPLE OF RULE CLASSIFICATION USING TUPLE SPACE.

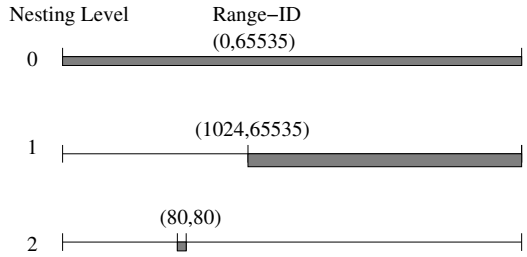

Fig. 2. Assigning values for ranges, based on the Nesting Level and the Range-ID.

the protocol fields, the value is simply a boolean: $'1'$ if a protocol is specified, $'0'$ if a wildcard is specified [8][9][10]. The number of specified bits in a port range is less straightforward to define. The authors introduced the concept of nesting levels and Range-IDs to define the tuple value for port ranges. The nesting level specifies the layer of the hierarchy and the Range-ID uniquely labels the range within its layer. In this manner, all port ranges can be converted to a (Nesting level, Range-ID) pair. We present an example to illustrate Range-IDs in the following. The full range, in this example (0-65535) always has the id 0. The two ranges at level 1 namely, (0, 1023) and (1024, 65535) in our example receive id 0, and 1, respectively. The example of mapping a port range to a nesting level and a Range-ID for Table II is depicted in Figure 2.
In the following, we illustrate how a search key is constructed from a packet based on a tuple. A search key for the tuple [8, 24, 2, 0, 1] is constructed by concatenating the first octet of the packet source address, the first three octets of the packet destination address, the Range-ID of the source port, the range at nesting level 2 covering the packet source port number, the Range-ID of the destination port range at nesting level 0 covering the packet destination port number, and the protocol field. Finally, all algorithms using the tuple space approach involve a search of the tuple space or a subset of the tuples in the tuple space.

IV. Analysis of Results

In this section, we present the analysis and discussion of results for the rule-set databases.

A. Results

In this section, we describe the analysis of rule-set databases. In the presented diagrams, all of observed values have been normalized to the number of rule-set entries in the related database. The distributions of source and destination IP address prefix bits are depicted in Figure 3.

From the Figure 3 (A), we can observe that most of source IP addresses in the rule-set databases have long prefix bits. The source IP addresses with 32 valid bits have the highest rank since most of the rules for the source IP addresses are applied with a specific IP address. Additionally, we can observe that some of rules do not utilize any bits in their source IP address fields which means these rules are employed to apply general policies.

Figure 3 (B) depicts the destination IP address prefix bits for different rule-set databases. From the Figure 3 (B), we can follow that most of the destination IP addresses have long prefix bits associated with them. The behavior of destination IP addresses is similar to source IP addresses in the rule-set databases.

The source port number distribution is depicted in Figure 4.

Figure 4 (A) depicts the distribution of source port numbers in the rule-set databases *fw1* and *acl1*. The lowest value for the source port distribution belongs to *fw1-1k*. Based on the observations, the distribution

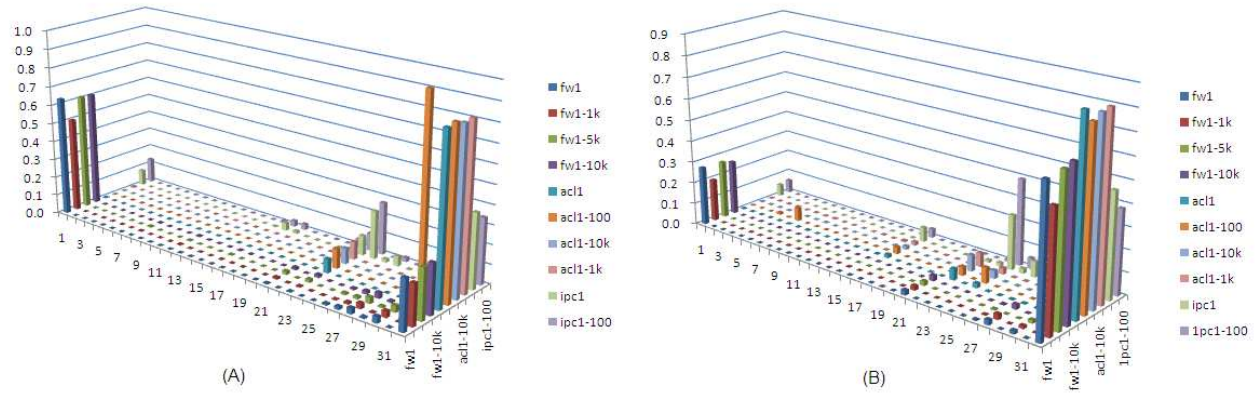

Fig. 3. IP address prefix bits distribution. (A) Valid bits of source IP addresses. (B) Valid bits of destination IP addresses.

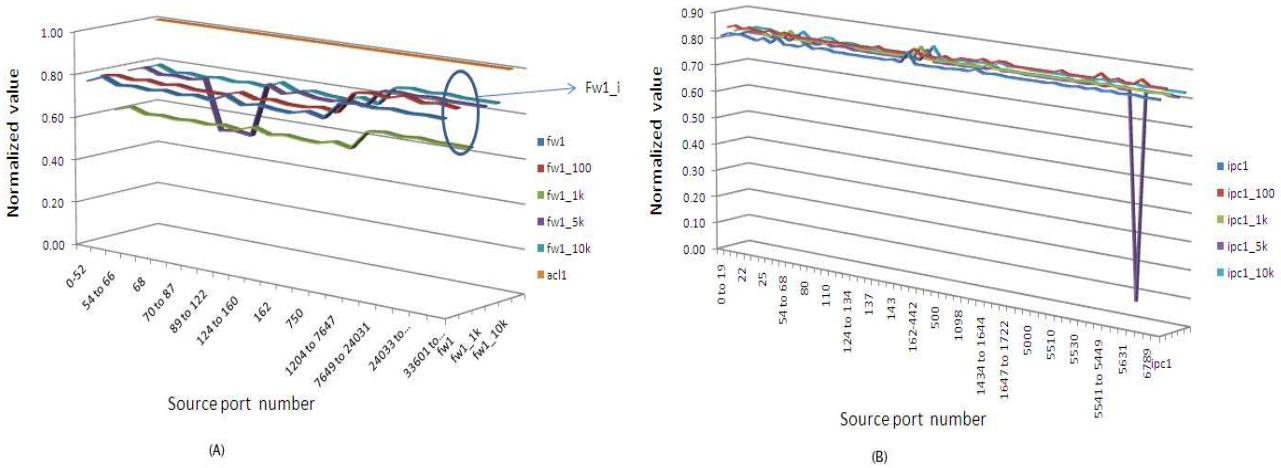

Fig. 4. Source port number distribution. (A) Source Port number distribution for *fw1*, *fw1-100*, *fw1-1k*, *fw1-5k*, *fw1-10k* and *acl* rule-set databases (the value of ports for *acl1-100*, *acl1-1k*, *acl1-5k* and *acl1-10k* are equal to *acl1*). (B) Source port distribution for *ipc1*, *ipc1-100*, *ipc1-1k*, *ipc1-5k* and *ipc1-10k* rule-set databases.

of source ports for rule-set databases *acl1* is the same in all of *acl* rule-set databases and is equal to number of rules in the *acl* rule-set database. It is due to the fact that *acl* rule-set database has been extracted from Access Control List in the firewalls and routers. All of them are equal to '1' and we only depicted the graph of *acl1* (real rule-set database). Figure 4 (B) is the same diagram for the rule-set database *ipc1*. From the figures 4 (A) and 4 (B), we can conclude that a wide range of different port numbers appear in most of the rules. The distribution of destination port numbers is depicted in Figure 5.

Based on the Figure 5, the rate of destination port number distribution is lower than the source port number counterpart. In most network devices, the rules are defined based on the particular destination port numbers, on the contrary, the source ports are mostly defined as a range. As a result, the diagrams show increased rate for the source port number distribution compared to the destination port number distribution.

The range-size distributions for the source/destination port numbers in different rule-set databases are depicted in Figure 6.

In Figure 6, range-size represents the distance between the starting and ending port numbers in the range. As an example, the range-size of (10-110) with the starting value of '10' and ending value of '110' in the source or destination ports will be '100'. Figure 6 (A) depicts range-size for different rule-set databases in the source port numbers. we can observe that there are many wide ranges in source port numbers. The range-size '1' represents just a single port. The range-sizes greater than '2' and less than '536' did not occur in any of the rule-set databases and range-sizes '1' and '2' are observed in some databases. From Figure 6, it is evident that most of the source port numbers are defined with wide ranges. Figure 6 (B) depicts, range-size for different rule-set databases in the destination

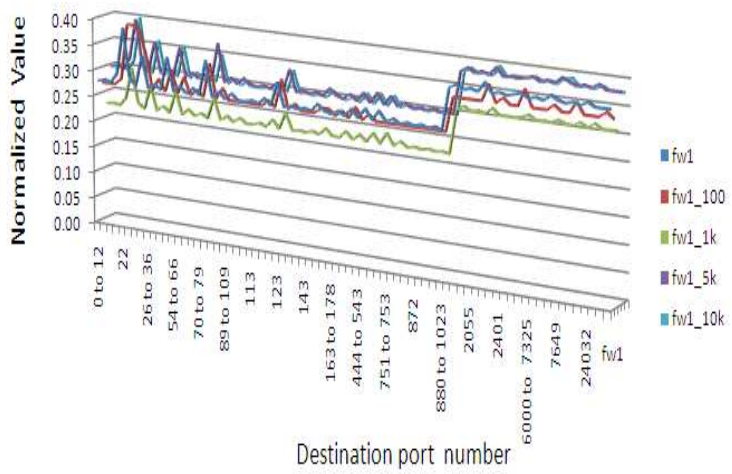

Fig. 5. Destination port number distribution for *fw1*, *fw1-100*, *fw1-1k*, *fw1-5k* and *fw1-10k* rule-set databases.

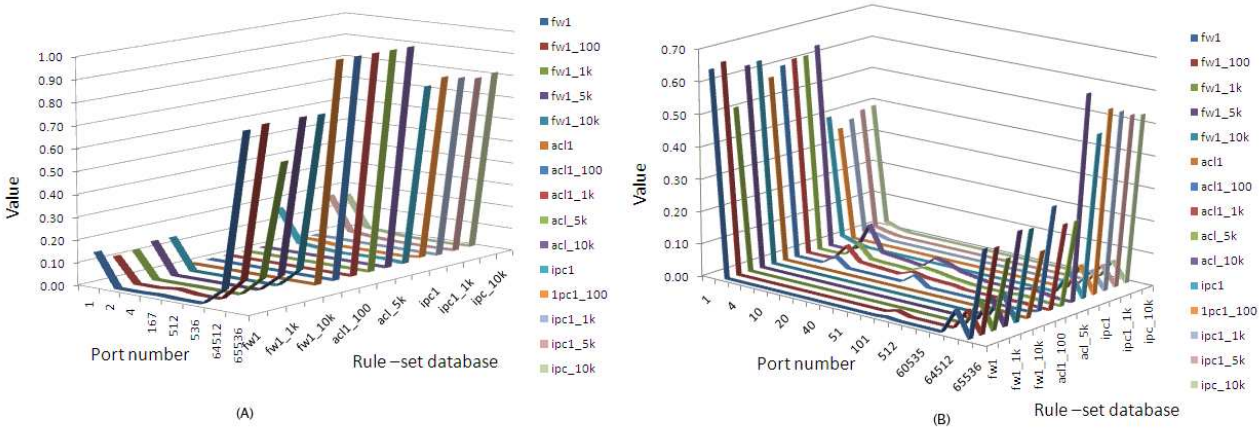

Fig. 6. Range-size distribution in different rule-set databases. (A) Range-size distribution for the source port numbers. (B) Range-size distribution for the destination port numbers.

port numbers. we can observe that most range-sizes belong to the very low value and high value destination port numbers. In the figure, range-size with value '1' appears more than other range-sizes. In other words, most of the rules are defined as specific port numbers and the remaining rules are generally defined with ranges of high values.

From Figures 6 (A) and 6 (B), we infer that in most of the rules, source port numbers are defined with wide ranges and most of the destination port numbers are specified as particular port numbers or as wide ranges.

B. Discussion

A well-known hash-based algorithm in packet classification is tuple space search algorithm that groups different rules using simple mapping and searches the IP packets in the tuples. In the mapping process the source and destination ports are mapped on the tuple space using Range-ID and nesting level concept. Utilization of the results of the rule-set databases' analysis can lead to the optimization of the hash-based algorithms. The results presented and the corresponding examinations show that the hash-based packet classification algorithms can be optimized using refined tuple space. In traditional tuple space, a search key is created using all of fields in the rules or packet headers and this consequently increases the number of tuples created as well as the search time needed for optimal matching. In an improved tuple space, for key construction, we take advantage of the range-size data to efficiently map the source and destination ports.

V. Overall Conclusions

In this paper, we presented an introduction to the well-known packet classification problem and stressed the contribution of rule-set database analy-

sis as a critical stage in packet classification algorithms. Processed rule-set data were demonstrated with illustrative diagrams regarding the distributions of source/destination IPs, source/destination ports and protocol fields. Based on the extracted information, researchers are expected to efficiently refine the existing or propose novel packet classification algorithms, which reveal higher performance and reduced processing time compared to traditional methods.

References

[1] A. Feldmann and S. Muthukrishnan, "Tradeoffs for packet classification," in *Proc. Int'l IEEE Conf. INFOCOM*, vol. 3, 2000, pp. 1193–1202.

[2] P. Gupta and N. McKeown, "Algorithms for packet classification," *J. IEEE Network*, vol. 15, no. 2, pp. 24–32, March-April 2001.

[3] T. V. Lakshman and D. Stiliadis, "High-speed policy-based packet forwarding using efficient multi-dimensional range matching," in *Proc. Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communication (ACM/SIGCOM)*, 1998, pp. 203–214.

[4] D. Taylor and J. Turner, "Classbench: A packet classification benchmark," Department of Computer Science And Engineering, Washington University in St. Louis, Tech. Rep. WUCSE2004-28, May 2004.

[5] P. Gupta and N. McKeown, "Packet classification on multiple fields," in *Proc. Int'l Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communication*, 1999, pp. 147–160.

[6] H. Song, "Evaluation of packet classification algorithms," 2006, http:// www.arl.wustl.edu/ hs1/PClassEval.html.

[7] V. Srinivasan, "A packet classification and filter management system," in *Proc. Int'l IEEE Conf. INFOCOM*, 2001, pp. 1464–1473.

[8] V. Srinivasan, S. Suri, and G. Varghese, "Packet classification using tuple space search," in *Proc. Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communication*, 1999, pp. 135–146.

[9] D. E. Taylor, "Models, algorithms, and architectures for scalable packet classification," August 2004.

[10] V. Srinivasan, "Ip lookup and packet classification," Ph.D. dissertation, Washnigton University, Saint lous, Missouri, August 1999.