

Kent Academic Repository

Full text document (pdf)

Citation for published version

Tripp, Gerald (2006) A parallel 'String Matching Engine' for use in high speed network intrusion detection systems. *Journal in Computer Virology*, 2 (1). pp. 21-34. ISSN 1772-9890.

DOI

<https://doi.org/10.1007/s11416-006-0010-4>

Link to record in KAR

<https://kar.kent.ac.uk/14483/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

A parallel 'String Matching Engine' for use in high speed network intrusion detection systems.

Gerald Tripp

University of Kent

About Author

Gerald Tripp is a Lecturer in Computer Science at the University of Kent.

Contact Details: The Computing Laboratory, University of Kent, Canterbury, Kent, CT2 7NF, UK, phone +44 1227 827566, fax +44 1227 762811, e-mail G.E.W.Tripp@kent.ac.uk

Keywords

Security, intrusion detection, finite state machine, string matching, high speed networks.

A parallel 'String Matching Engine' for use in high speed network intrusion detection systems.

Abstract

This paper describes a finite state machine approach to string matching for an intrusion detection system. To obtain high performance, we typically need to be able to operate on input data that is several bytes wide. However, finite state machine designs become more complex when operating on large input data words, partly because of needing to match the starts and ends of a string that may occur part way through an input data word.

Here we use finite state machines that each operate on only a single byte wide data input. We then provide a separate finite state machine for each byte wide data path from a multi-byte wide input data word. By splitting the search strings into multiple interleaved substrings and by combining the outputs from the individual finite state machines in an appropriate way we can perform string matching in parallel across multiple finite state machines.

A hardware design for a parallel string matching engine has been generated, built for implementation in a Xilinx Field Programmable Gate Array and tested by simulation. The design is capable of operating at a search rate of 4.7 Gbps with a 32-bit input word size.

Introduction

Network intrusion detection consists of monitoring computer networks for various types of security attack. This can be network wide monitoring (network based) or it can be at each individual host computer in the system (host based). Basic network security is provided by network firewalls, which act as an intermediary between the Internet and a local network – these filter network traffic on the basis of header fields in the packets such as the source and destination IP address and TCP port numbers. This type of filtering is good at blocking a large proportion of unwanted incoming traffic. However, some network attacks may be targeted at machines such as web and mail servers that need to be visible through the firewall. In this case, it may be necessary to look inside each incoming data packet to determine whether it represents a potential threat. We may then wish to block that traffic (intrusion prevention) or be able to generate an alert that potentially malicious traffic is present (intrusion detection). The problem is that we may have no particular field to examine inside the packet, and may need to search the entire packet. This is the standard technique that we use for intrusion detection: we first look at the header fields of the packet to see if the packet is potentially of interest and if so we then search the content of the packet for one or more related intrusion detection 'signatures'. These signatures are short search strings which are chosen as representing a high probability of an attack occurring when present, whilst having a low probability of occurring otherwise.

A lot of current intrusion detection systems are software based, the most well known example probably being Snort (Roesch, 1999). Software solutions can however have problems when presented with a high network load. One solution can be to use host based intrusion detection and to require each computer to perform its own intrusion detection. This however can be targeted by denial of service attacks to put the intrusion detection software on individual machines under heavy load. Host based solutions are also only possible if we are able to add intrusion detection software to each host system, and this may not be the case with some embedded systems.

Summary of this paper

This paper looks at the string matching part of intrusion detection and describes how it is possible to build a 'string matching engine' for implementation in a Field Programmable Gate Array (FPGA) that uses fine grained parallelism to improve its search rate. The method used is to operate on a multi byte input data word and to partition the matching operation between a set of Finite State Machines (FSMs), each of which processes one of the byte streams from a multi-byte wide network input and looks for parts of the search string. The results from these multiple FSMs are then combined in a particular way so as to determine whether a string has been matched across all the FSMs.

The next section describes the background and outlines some of the related work in this field. The following section describes the operation of the parallel string matching system proposed in this paper. The software section gives the results of processing multiple search strings and the resource requirements for various string set sizes and implementation options. The next section gives details of a hardware design for a string matching engine and its performance and resource requirements. The final section gives conclusions and ideas for further work.

Discussion

A lot of existing intrusion detection systems are software based, the most well known example probably being Snort (Roesch, 1999). Many improvements have been made to Snort by optimising the order in which data is compared. Work by (Kruegel & Toth, 2003) uses rule clustering and is implemented as a modified snort rule engine. This uses decision trees to reduce the number of comparisons made against incoming network data and uses a multiple string matching algorithm based on the work by (Fisk & Varghese, 2001).

A paper by (Abbes, Bouhoula & Rusinowitch, 2004) describes a system using a decision tree in conjunction with protocol analysis. The protocol analysis uses a specification file for the protocol being monitored and performs 'Aho-Corasick' (Aho & Corasick, 1975) string matching on only the appropriate parts of the data stream. This technique reduces the overall workload and also reduces the number of false positives as compared with performing matching on the entire data packet or using simple offset and depth constraints.

Work by (Paul, 2004) looks at distributed firewalls and implements stateful packet classification spread across consecutive firewalls. This helps to spread the workload between separate machines.

It can be difficult to perform intrusion detection in software at high network traffic rates and hardware solutions may be required. Software solutions being essentially sequential also suffer from performance problems as we increase the number of rules; (Cho & Mangione-Smith, 2004) state that a software system with 500 rules may have difficulty in sustaining a throughput of 100 Mbps. Hardware solutions have different limitations; we can often increase the number of rules without affecting throughput because of the use of parallelism – the cost of increasing the number of rules may be an increase in hardware resource utilisation instead.

Overview of existing solutions

A number of hardware based string matching systems for intrusion detection have been described in the literature; an overview of some of the techniques is given below.

A product called ClassiPi from PMC-Sierra is described by (Iyer, Kompella & Shelat, 2001), this is a classification engine and implemented as an application specific integrated circuit (ASIC). This

device allows software-like algorithms to be used for various packet classification and packet inspection operations, including the use of regular expressions to search the contents of packets.

Work by (Attig & Lockwood, 2005) uses Bloom filters to perform string searching. Bloom filters provide an efficient method to perform searching for a large number of strings in parallel, but suffer from the disadvantage of producing false positive matches. Attig and Lockwood show that Bloom filters can be used as a very efficient front end to remove the bulk of the network traffic that is known to be benign before input into a conventional software intrusion detection system.

(Cho et.al., 2004) describe a system that uses multiple matching systems, each of which will search incoming network data for a set of distinct string 'prefixes'. For each possible string prefix, their system will lookup the remaining part of the string that must be compared sequentially against the incoming data to determine whether that string is actually present. Multiple strings with identical prefixes need to be distributed between different matching systems.

An interesting approach is taken by (Baker & Prasanna, 2004), who have a series of input comparators for each data byte of interest – the output of these comparators each feed into a pipeline of flip-flops. Strings can be identified by the use of an AND function that looks for all the required data bytes for a string in the appropriate positions within the pipeline. They show that this can be extended to operate with multi-byte input data by the use of multiple sets of pipelines and looking for strings across the set of pipelines at all byte alignments.

Finite state machine approaches

A number of systems have been designed that use Finite State Machines (FSM) to perform the searching – most of these use a Deterministic Finite Automata (DFA) to implement string matching. This type of FSM has sets of states, inputs and outputs; the FSM can be in one of its states and there is a mapping between each pair of current state and input to the next state and output. When used in string matching, we use the FSM state to define how much of a string we have matched so far.

The approaches taken by (Sugawara, Inaba & Hiraki, 2004) and by (Tripp, 2005) is to first compress multi-byte input data into a number of different patterns that are of interest and then to use DFAs to perform string matching several bytes at a time. (Moscola, Lockwood, Loui & Pachos, 2003) convert regular expressions into DFA that operate one byte at a time and show that this can be used to perform matching for standard spam-assassin rules without creating too many DFA states.

A different approach is taken by (Franklin, Carver & Hutchings, 2002), who implement Non-deterministic Finite Automata (NFA) in hardware to perform matching of strings from the Snort rule set, this approach first being proposed by (Sidhu & Prasanna, 2001). This was extended by (Clark & Schimmel, 2004) to operate with multi byte input data.

The text by (Hopcroft, Motwani & Ullman, 2001) gives a comprehensive coverage of Deterministic and Non-deterministic Finite Automata.

String matching algorithms

There are many string matching algorithms described in the literature, most of which were originally devised for software implementation. A hardware implementation has slightly different requirements than that for a software implementation and may well need to be less complex. For efficiency it is more common to build systems that work on a stream of data, rather than providing random access to the contents of a buffer; ideally we would like the string matching to operate at a deterministic rate to avoid the need for buffering.

The fastest method of matching strings is considered to be the Boyer-Moore algorithm (Boyer & Moore, 1977) and its successors. This performs string matching on a 'right to left' basis and skips forward on a mismatch. This gives an average performance that is usually sub-linear, but a worst case performance that may require us to look at some input bytes many times.

The 'Knuth Morris Pratt' (KMP) algorithm (Knuth, Morris & Pratt, 1977), performs matching on a left to right basis and on mismatch will use the longest partial match as a starting point for further matching. The algorithm can be adapted to operate at deterministic data rate and not re-examine input data on a mismatch.

The Aho-Corasick algorithm (Aho et.al., 1975) matches several strings at the same time. This works by constructing a trie containing the various strings and this is traversed as the data arrives. As with KMP, this can also be modified to operate at a deterministic rate only looking at each input data item once.

Both KMP and Aho-Corasick can be implemented by creating a FSM that operates at one input data item per clock cycle and are therefore ideal for hardware implementation. A common method of implementation for both these algorithms uses a maximum FSM size of an initial state and one state per search character (in one or all strings). When using Aho-Corasick, we would have fewer states when common prefixes of search strings enable us to share a FSM state. The state transition information in both cases will vary in complexity determined by whether on mismatch of a partly matched string there exists a suffix of the data matched that forms a smaller partial match of that string (or another).

Parallel string matching

From the work presented by (Sugawara et.al., 2004) and (Tripp, 2005), we can see that high performance can be obtained by creating a FSM that will match multiple bytes in the same clock cycle. However this has the overhead of compressing the input data so as to present a small input word to the FSM. A second issue is that the start and ends of strings have a high chance of appearing part way through an input data word, so we may need to match parts of the start and end of a string with 'wild card' characters.

It is far easier to match data from an 8-bit input bus, but this does not give such good throughput. The solution proposed here is to use multiple finite state machines in parallel to process the input data. Course grained parallel FSM solutions have already been implemented, such as the work described by (Moscola et.al., 2003), where input packets are allocated to a number of content scanners on a round robin basis. We propose a fine-grained form of parallelism, where multiple finite state machines process each packet in parallel.

Parallel finite state machines

The approach we take here is to provide a finite state machine for each byte stream from a multi-byte input data word. If we have a w -byte wide input word, then we can use w separate finite state machines, each of which are looking for all w instances of the 'substrings' made up from a w -way interleave from the search string. An example of such a system is shown in Figure 1.

A related, but different, approach is taken by (Tan & Sherwood, 2005) who use multiple FSMs running in parallel to match a sequence of bits, with each FSM matching a particular bit position from the input data.

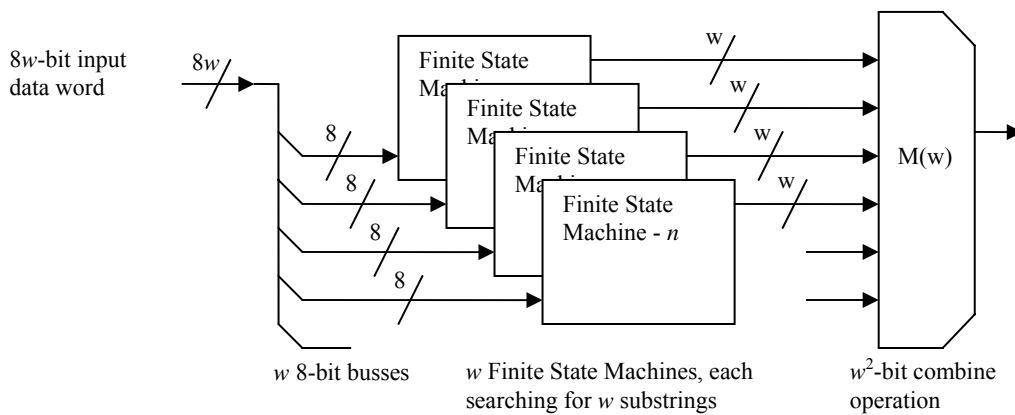


Figure 1 - Matching Interleaved substrings.

All w instances of our FSM are identical, and each will be looking for all w substrings. Each FSM has a w -bit Boolean ‘match vector’ output to specify the sub-strings matched in any clock cycle. If we find all w substrings appearing in an appropriate order across all w finite state machines at the correct time, then we will have found our search string. We can see an example of a set of substrings of a given search string when $w = 4$ in Figure 2.

Word size = 4			
Search string =	"the-cat-sat-on-the-mat"		
Substring 0 =	" e t t - - "	=	"ett--"
Substring 1 =	" - - - t m "	=	"---tm"
Substring 2 =	"t c s o h a "	=	"tcsoha"
Substring 3 =	" h a a n e t "	=	"haanet"
(The substrings are sorted by the order of completion, the reason for which will be explained below.)			

Figure 2 - Interleaved substrings.

By sorting our substrings on the basis of the order of completion of the match, we have a sequence in byte terms of w consecutive substring matches. However, we are processing our data on the basis of a w -byte input word. The string may be aligned in one of w different ways, with the last w bytes occurring in one or two input data words – the occurrence of each of the last w bytes of the search string relate to the instant when each of the related substring matches will occur. We define here an alignment of c as meaning that of the last w bytes of the search string, c of these will occur in one input word, followed by $(w - c)$ in the following input word, where: $0 \leq c < w$.

Byte stream x is being monitored by finite state machine x . Each of the finite state machines is searching for all w substrings, and has a Boolean ‘match’ output for each substring y . Thus we have a group of w^2 FSM outputs: $O_{x,y}$, where $0 \leq x < w$ and $0 \leq y < w$, relating to whether FSM x has detected substring y in the current clock cycle. We are also interested in whether string matches occurred in the previous clock cycle, and $O'_{x,y}$ is a delayed (pipelined) copy of $O_{x,y}$ from the previous clock cycle.

Taking the case where $w = 4$ and $c = 1$ for the string in Figure 2, we have the alignment shown in Table 1.

		Input word					
		n-5	n-4	n-3	n-2	n-1	n
Input Byte	0		-	-	-	t	m * ₁
	1	t	c	s	o	h	a * ₂
	2	h	a	a	n	e	t * ₃
	3	e	t	t	-	- * ₀	

*_s indicates when a match occurs for substring S.

Table 1 - String match at alignment c=1.

We define $M_c(w)$ as being a Boolean operation specifying whether a match occurs at alignment c , in a system with a word size w . In our example above, we have $c = 1$ and $w = 4$; we can see from Table 1, that $M_1(4)$ is as shown in Equation 1.

$$M_1(4) = O'_{30} \cdot O_{01} \cdot O_{12} \cdot O_{23}$$

Equation 1 - Match occurs at alignment c=1, for word size w=4.

This follows a very simple pattern, and we can produce a general formula for $M_c(w)$. Our complete string match is then defined as $M(w)$ which determines whether the match occurs at any of the w possible alignments. This is shown in Equation 2.

$$M_c(w) = \bigwedge_{i=0}^{w-1} \text{if}(i \geq c) \text{ then } (O_{(i-c)i}) \text{ else } (O'_{(i+w-c)i})$$

$$M(w) = \bigvee_{c=0}^{w-1} M_c(w) = \bigvee_{c=0}^{w-1} \left(\bigwedge_{i=0}^{w-1} \text{if}(i \geq c) \text{ then } (O_{(i-c)i}) \text{ else } (O'_{(i+w-c)i}) \right)$$

Equation 2 – The Combine operation.

(Note that in Equation 2, we use \bigvee and \bigwedge to represent the Boolean ‘inclusive-or summation’ and ‘and product’ respectively.)

The combine operation $M(w)$ is independent of the search string and can be implemented as a

fixed logic function for a given value of w . We also need $\sum_{x=1}^{w-1} x$ D-type flip flops to generate the

delayed versions of some of the inputs. As an example, the combine operation required for a system with a word size of 4 bytes is shown in Equation 3.

$$M(4) = O_{00} \cdot O_{11} \cdot O_{22} \cdot O_{33} + O'_{30} \cdot O_{01} \cdot O_{12} \cdot O_{23} + O'_{20} \cdot O'_{31} \cdot O_{02} \cdot O_{13} + O'_{10} \cdot O'_{21} \cdot O'_{32} \cdot O_{03}$$

Equation 3 – Combine operation for a 4-byte word.

This requires four 4-input ‘and’ gates, one 4-input ‘or’ gate and six D-type Flip-flops.

In terms of overall complexity, the move from a standard byte-wide Aho-Corasick multi-string matching system to the technique described here requires us to replace a single FSM with w instances of a new FSM for matching sub-strings and one instance of the combine operation described above. The new FSM will have a similar number of states to the original, but will require a factor of w increase in the number of match outputs. Actual resource utilisation will depend on many parameters relating to the FSM implementation as will be shown later. The resources required for the combine operation are trivial for small values of w – but will grow rapidly in size with w as it implements a w^2 input Boolean function.

Implementation

Each FSM has to be able to match multiple substrings, and this can be done using the Aho-Corasick multiple string matching algorithm. As we are using a multiple string matching algorithm we can actually use each FSM to search for the substrings for several different search strings.

The method used here for the FSM implementation is table based – the reason for taking this approach is that we are able to have a fixed core of logic for any FSM (of a given size) and we determine the operation performed by the FSM by specifying the contents of the FSM table. The state transition table for such a FSM is very redundant, and this can be implemented using the type of FSM implementation described by (Sugawara et.al., 2004), as shown in Figure 3.

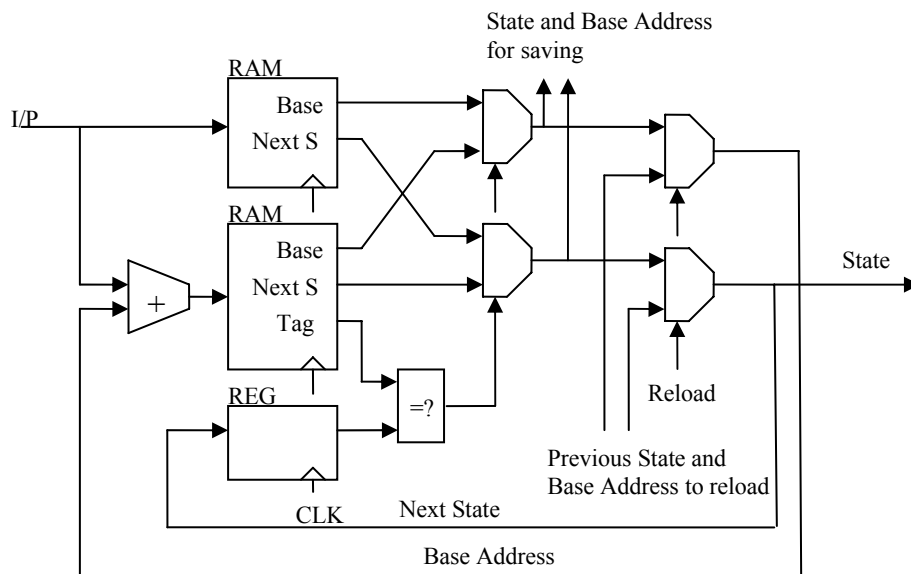


Figure 3 - Finite State Machine Implementation from (Sugawara et.al., 2004).

The algorithm by (Sugawara et.al., 2004) works on the basis that for a given input value i , a large proportion of transition table entries for current state s will be the same as for the IDLE (or initial) state. The algorithm uses a default table that contains table entries for all input values of i in the IDLE state. All we need in addition to this are the entries from the full state transition table that differ to the entries in the default table – this difference table is typically very sparse.

Table 2 gives an example of a simple FSM that searches for the single string “abcabc”. The input is a numerical value that relates to the character shown in brackets. The state represents the portion of the search string that has been matched. The tables contain the next state for the FSM and in this example the match succeeds when the FSM is in state 6.

Original State Transition Table		Default Table	Difference Table												
Current State		Current State													
	:		:												
	'a'		'a'												
	'ab'		'ab'												
	'abc'		'abc'												
	'abca'		'abca'												
	'abcab'		'abcab'												
	'abcabc'		'abcabc'												
	0		0												
	1		1												
	2		2												
	3		3												
	4		4												
	5		5												
	6		6												
Input	0(a)	1	1	1	4	1	1	1							
	1(b)	0	2	0	0	5	0	0			2			5	
	2(c)	0	0	3	0	0	6	0				3			6
	3(x)	0	0	0	0	0	0	0							

Table 2 – Simple FSM to match search string ‘abcabc’.

To find the next state for any current state and input, we first look in the difference table. If this does not have an entry then we use the value from the default table instead.

This difference table is decomposed into a series of state vectors, and these are packed together (overlapping) into a one-dimensional packed array – carefully avoiding any collisions between active entries. Each entry in the packed array is tagged with the current state it belongs to. To retrieve an entry from the packed array we need to know the base address of the state vector for the current state in the packed array and then use the current input as an offset from that point. If the entry fetched from the array has a tag that is equal to the current state, then we have found a valid difference table entry – if not, there is no entry for the current state and input in the packed array, so we use the value from the default array for the current input.

We can see an example in Table 3 of the how the FSM in Table 2 is converted into this format. To improve performance, each entry (in both arrays) also contains the base address of the state vector in the packed array for the next state (NS).

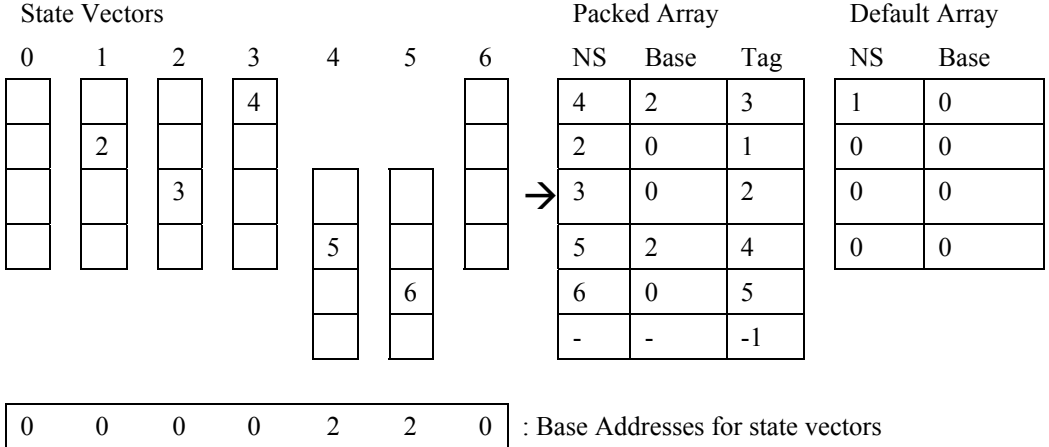


Table 3 – Packed and Default arrays, using the method described by (Sugawara et.al., 2004).

The algorithm by (Sugawara et.al., 2004) gives a significant memory saving for large FSMs, as we avoid the use of the two dimensional arrays. It is difficult to a give a figure for the resource utilisation based on string length or number of strings, as the memory required will be determined by how well the state vectors can be fitted together into the packed array.

Modifications to the FSM implementation

For our work, the design in Figure 3 was setup to operate on one packet at a time, by providing a restore input corresponding to the IDLE state. After testing the FSM design, it was found that performance was limited by the adder carry chain used in the implementation of the "+" operation that provides the index into the packed array. On investigation, it was found that it was possible to replace the "+" operator with a "bitwise XOR" operation. This is not as efficient as it will constrain any state vector to be within a single 2^n sized block of memory for a data input bus size of n-bits – in practice however it is not found to cause much reduction in efficiency, as seen later. The advantage is that it makes each bit in the address calculation independent which improves the hardware performance.

The state from each FSM is fed into a state decoder table that generates a substring match vector specifying all of the substrings that complete matching in the current state. The use of a table for this operation avoids needing to build specific logic for each string set. As strings shorter than the word size can generate a match from a subset of the FSMs, we allow one or more of its substrings to be the 'null string' which will match in any state, thus we always require a match from all FSMs irrespective of search string length.

Rule processing

Rather than generating a specific piece of hardware for a given rule set, it was decided that we should identify an efficient size of 'string matching engine' and then instantiate a number of these to cover the set of strings. We will not know in advance how many strings will fit into a FSM of any particular size, as this will depend on how compact the packed array can be made. The best size of FSM will depend on a number of factors, but will relate in particular to the memory resources available in the hardware. As we don't know in advance how many strings we can fit into a given FSM, we need to take an iterative approach and try increasing numbers of strings to see how many will actually fit.

Rule sets such as those defined by snort will allow us to have content matching that is case independent. We can deal with this by allocating these strings to separate 'string matching engines' to the ones used for strings that are case dependent and pre-pending an input function that maps all upper case letter to lower case.

Software

Software was written to take a set of strings and to build an Aho-Corasick trie for performing the matching. The design was optimised using standard techniques to enable the matching to be performed at a rate of one byte per clock cycle. From this, a state transition table was produced and then compressed using the technique described by (Sugawara et.al., 2004) and outlined above.

The first stage was to choose a sensible size for the FSM. The software was modified so that instead of reading in all the search strings, it stops after a certain limit of search characters had been exceeded and the memory resources required for that amount of search characters reported. This was repeated for a range of maximum numbers of search characters – the search strings being taken from a randomised order set of case dependant rules from the hogwash (Larsen & Haile, 2001) "insane" rule set. The operator for the packed array index was chosen as the ADD operator.

The tests were performed for a range of input bus sizes, and the memory requirements for a single 8-bit slice are given in Figure 4. We see that the amount of memory required increases with the input bus width, as the number of substrings increases with the number of 8-bit slices.




Figure 4 - Memory use.

The traces in Figure 4 are roughly linear for a range of total characters from 200 to 2000. The approximate memory requirements B_w in bits for word size w and character count s in this range are shown in Equation 4. Memory usage will of course vary with the particular set of strings chosen.

$$B_8 = 0.10s + 51, \quad B_{16} = 0.19s + 53, \quad B_{32} = 0.33s + 33, \quad B_{64} = 0.35s + 54$$

Equation 4 – Approximate memory use for a range of 200 to 2000 search characters.

When we get to a word size of 64-bits, the total amount of memory required does not increase as much as expected, as there are an increasing number of short identical substrings, including null strings. Calculations here show the exact amount of memory required – in practice the memory will only be available in particular sizes, as shown later.

Interestingly, the amount of memory required per search character increases with the total amount of characters in the search strings. This is partly due to the memory requirements of the state decoder, but this effect is present even if we don't take this into account. We would expect to get some gain as we increase the number of search strings as we should have nodes within the trie shared between multiple search strings. We can see this effect in Figure 5.




Figure 5 - FSM resources.

The number of trie nodes per search character decreases with the total number of search characters as expected; however this effect is counteracted by the increase in the number of ‘non-default edges’. The ‘non-default edges’ are transitions from one node to another that are recorded in the packed array and this relates to the FSM becoming more complex and there being more interconnectivity between nodes. The overall effect is that the number of non-default edges per search character increases with the number of characters. The total memory requirements for the packed array are roughly proportional to the number of non-default edges, as each ‘non-default edge’ will require its own entry in the packed array.

All three traces in Figure 5 are roughly linear up to a total of 1700 characters in the search strings. The approximate results in this range, for a total of s characters are shown in Equation 5.

$$nodes / char = 0.79 - 9 \times 10^{-5} s$$

$$edges / node = 3.3 \times 10^{-3} s + 0.82$$

$$edges / char = 2.0 \times 10^{-3} s + 0.81$$

Equation 5 - Approximate resource usage for up to 1700 search characters.

From Figure 5, it can be seen that when using this style of implementation it is not necessarily the best option to have large Aho-Corasick FSMs. The resources used appear to be lower when only a small number of strings are searched for; this however will be dependent on the sizes of memory available for the various FSM tables.

Determining an optimal FSM size

The software was modified to re-run a number of tests for a fixed input word size of 32-bit, using variations of the algorithms. The tests were run for an increasing number of search strings and the total memory resources required were calculated for implementation within a Xilinx Virtex-II FPGA (“Xilinx Virtex-II”, 2005) – this type of FPGA contains 18Kbit Block RAM primitives (BRAMs). As an experiment, we also test the effect of preceding each FSM input with a custom built compression table to reduce the redundancy in the input data – as shown in Figure 6.




Figure 6 - One 8-bit slice of matching system.

The tests were run with either raw or compressed input data and with either ADD or XOR used for packed array indexing. The results are shown in the left graph of Figure 7; the four traces are very close together, and an enlarged section is given (for clarity) in the right graph where the resource utilisation is the lowest.

The use of compression did not have much effect when we use a large number of strings, however with a small number of strings the memory used increased because the extra memory needed for input compression was greater than the memory saved within the FSM tables. The choice of ADD or XOR algorithms had very little effect.

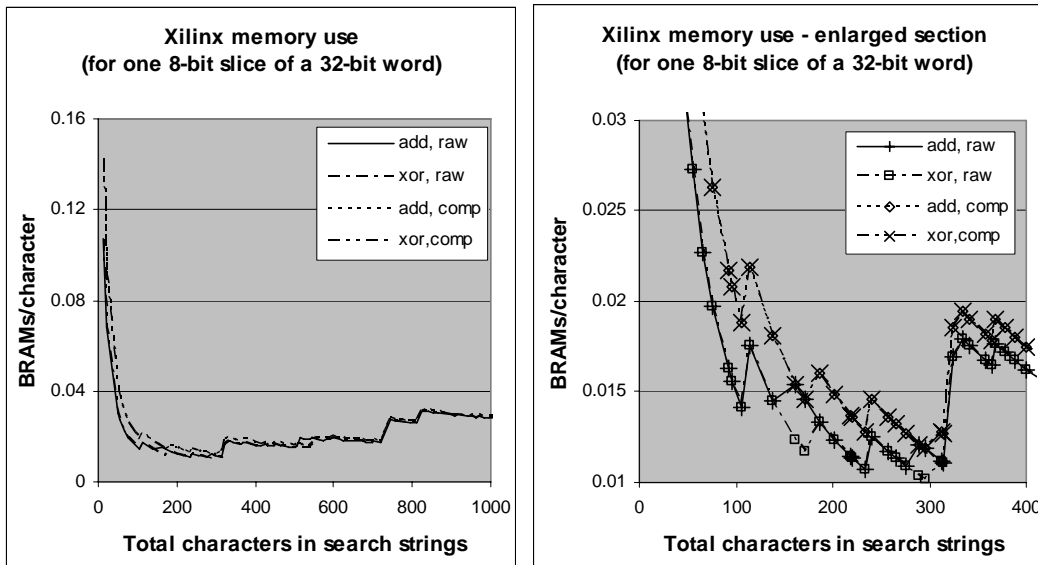


Figure 7 - Xilinx memory use vs. the number of search characters.

From the results shown in Figure 7, the best option appears to be a small FSM implementation dealing with a maximum of 200-300 search characters, and which uses 3 BRAM components. The best results for each of the four algorithm combinations are shown in Table 4.

FSM Input	Packing Algorithm	Search Strings	Search Characters
Raw	ADD	22	275
Raw	XOR	24	295
Compressed	ADD	18	234
Compressed	XOR	18	234

Table 4 - Maximum number of search strings for a 3 BRAM implementation.

Hardware implementation

A VHDL model was built of a 32-bit string matching engine that consisted of four 8-bit wide matching ‘slices’ and a unit to combine together the results – as shown in Figure 8. On the basis of the results above, a decision was made not to use input compression and to use the XOR function for indexing into the packed array for the FSM. The VHDL model was tested by simulation, the design synthesised and built for a Xilinx XC2V250-6 FPGA to determine its performance and resource utilisation. The design was also simulated “post place and route” to test the resulting FPGA design.

The parameters of the FSM design were taken from the rule processing results of the previous section. Each FSM has an 8-bit input, an 8-bit state variable and a 108-bit substring match output. (Note: The value of 108 was chosen as it is a multiple of one of the BRAM memory widths, which is 36-bits.) Four instances of the FSMs were used with a fixed combine operation to generate a matching engine having a 32-bit data input and a 27-bit match output. This is capable of matching up to 27 search strings in parallel, depending on the length of the strings. The use of the match vector output enables us to indicate matches of multiple search strings occurring at the same time; this match vector output could be used to generate an indication of which strings occurred within a given input data packet (including the detection of multiple matches of different strings).

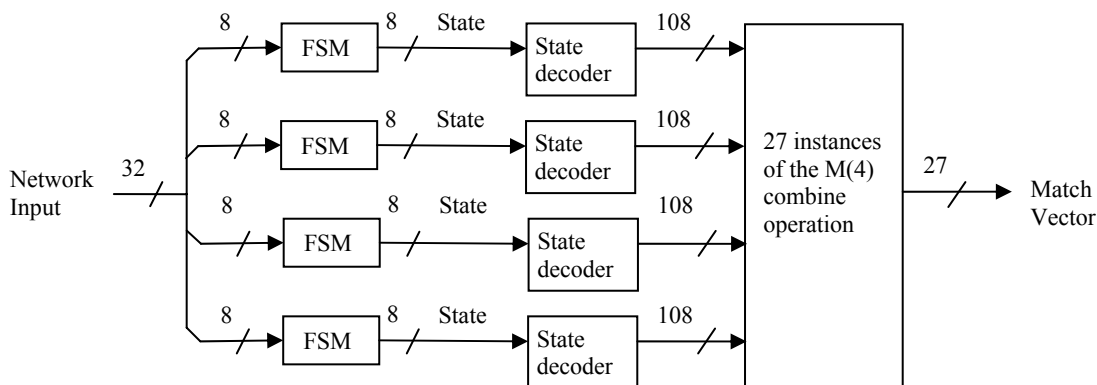


Figure 8 - Matching Engine.

Performance and resource utilisation

The VHDL model was first configured for using a bitwise XOR operator for the FSM table index operation and this gave a minimum clock cycle time of 6.7ns (149 MHz) – given the 32-bit input, this corresponds to a search rate of around 4.7 Gbps.

The resources required for a Xilinx XC2V250-6 FPGA were as follows:

- 12 Block RAM components (out of a total of 24)
- 250 logic slices (out of a total of 1536)

We can see from the above, that the size of any design will be limited by the Block RAM resources. The FPGA component used as the target of these experiments is however by current standards rather small. Taking the top of the range Virtex4FX FPGA as a comparison, we should be able to fit 46 of these matching engines within the FPGA, using all of the BRAM resources and around 20% of the logic slices. This should enable us to perform a parallel search of around 900 search strings.

For comparison, the VHDL model was rebuilt to use an ADD operator for the packed array indexing and this gave a minimum clock cycle time of 8.4 ns (119 MHz – giving a 3.6 Gbps search rate). This confirms the earlier assertion that using the bitwise XOR operator for the packed array indexing would be faster than using ADD.

Testing

The design has been tested by simulation with a large set of artificial input data containing various combinations of the strings being searched for, including: isolated instances of search strings; combinations of search strings at various spacing and overlap; and some strings rearranged in all 24 variations of the byte ordering in a group of 4 bytes. These tests were repeated for all four byte alignments of the input data – giving a total of 288 test cases. The results of the tests were compared with the expected outcomes to ensure that the search strings only matched as and when was expected. All tests passed correctly both for the original VHDL design and the post place and route simulations.

Pattern	Byte position that match occurs		
	"cybercop"	"gOrave"	"login: root"
----cybercop=====	12		
----ycebcprp=====			
----ybcecorp=====			
----cybercybercop=====	16		
----gOrave=====		9	
----login: root=====			14
----logOrave=====		11	
----killogin: root=====			17

Table 5 - A few of the test patterns used and the expected results.

A few examples of patterns used to test matching and the expected results are shown in Table 5.

Conclusion

This paper describes the design and simulation of a parallel algorithm for the implementation of high speed string matching; this uses fine-grained parallelism and performs matching of a search string by splitting the string into a set of interleaved substrings and then matching all of the substrings simultaneously.

We show that the FSM implementation technique described by (Sugawara et.al., 2004) can be modified by the use of bitwise XOR in place of ADD for the indexing operation to improve its performance. We also see that this implementation can be optimised in terms of resource utilisation by the choice of FSM size.

A VHDL model of a string matching engine based on the above ideas has been produced, synthesised and built for a Xilinx FPGA and tested via simulation. The results show a search rate of around 4.7 Gbps for a 32-bit input word. The design is table based and changes to the search strings can be made by generating new contents for the tables rather than having to generate a new logic design – this is particularly important for systems being updated in the field.

Future Work

One area where the resources in this design could be reduced is in the state decoder table – which accounts for 50% of the memory resources. This gives a substring match vector for the current state of the FSM – thus showing which substrings match in a given state. This table could be replaced with a piece of logic, but this would need to be rebuilt for every set of strings.

Further work is needed to see if the memory requirements for the state decoder can be decreased, possibly taking advantage of the redundancy that exists within this table. This could for example be replaced by a two stage decoder design. Finally it would also be interesting to see if any parts of the state decoder could be implemented as fixed logic.

References

- Abbes, T., Bouhoula, A., & Rusinowitch, M. (2004). Protocol Analysis in Intrusion Detection Using Decision Tree. In proceedings of International Conference on Information Technology: Coding and Computing (ITCC'04), Volume 1 (pp. 404-408). Las Vegas, Nevada.
- Aho, A.V., & Corasick, M.J. (1975). Efficient string matching: an aid to bibliographic search. *Communications of the ACM*, 18(6), 333-340.
- Attig, M., & Lockwood, J.W. (2005). SIFT: Snort Intrusion Filter for TCP. In Proceedings of IEEE Symposium on High Performance Interconnects (Hot Interconnects-13). Stanford, California.
- Baker, Z.K., & Prasanna, V.K. (2004). A methodology for Synthesis of Efficient Intrusion Detection Systems on FPGAs. In proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines FCCM '04. Napa, California.
- Boyer, R.S., & Moore, J.S. (1977). A Fast String Searching Algorithm. *Communications of the Association for Computing Machinery*, 20(10), 762-772.
- Cho, Y., & Mangione-Smith, W. (2004). Deep Packet Filter with Dedicated Logic and Read Only Memories. In proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines FCCM '04. Napa, California.
- Clark, C., & Schimmel, D. (2004). Scalable Multi-Pattern Matching on High-Speed Networks. In proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines FCCM '04. Napa, California.
- Fisk, M., & Varghese, G. (2001). An Analysis of Fast String Matching Applied to Content-Based Forwarding and Intrusion Detection (successor to UCSD TR CS2001-0670, UC San Diego, 2001). Retrieved 9 March 2006, from <http://public.lanl.gov/mfisk/papers/setmatch-raid.pdf>
- Franklin, R., Carver, D., & Hutchings, B.L. (2002). Assisting Network Intrusion Detection with Reconfigurable Hardware. In proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines FCCM '02 (pp.111-120). Napa, California, USA.
- Hopcroft, J.E., Motwani, R., & Ullman, J.D. (2001). *Introduction to Automata Theory, Languages and Computation*, 2nd Edition : Addison Wesley.
- Iyer, S., Rao Kompella, R., Shelat, A. (2001). ClassiPi: An Architecture for fast and flexible Packet Classification. *IEEE Network*, 15(2), 33-41.
- Knuth, D.E., Morris J.H., & Pratt, V.B. (1977). Fast pattern matching in strings. *SIAM Journal of Computing*, 6(2), 323-350.
- Kruegel, C., & Toth, T. (2003). Using Decision Trees to Improve Signature-based Intrusion Detection. In Proceedings of the 6th Symposium on Recent Advances in Intrusion Detection (RAID2003), Lecture Notes in Computer Science, LNCS 2820 (pp. 173-191). Springer Verlag.
- Larsen, J., & Haile, J. (2001). Securing an Unpatchable Webserver ... HogWash. Retrieved 9 March 2006, from <http://www.securityfocus.com/infocus/1208>

- Moscola, J., Lockwood, J., Loui, R.P., & Pachos, M. (2003). Implementation of a content-scanning module for an internet firewall. In proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines FCCM '03. Napa, California.
- Paul, O. (2004). Improving Distributed Firewalls Performance through Vertical Load Balancing. In proceedings of Third IFIP-TC6 Networking Conference, NETWORKING 2004, Lecture Notes in Computer Science, LNCS 3042 (pp. 25-37). Springer-Verlag.
- Roesch, M. (1999). Snort - Lightweight Intrusion Detection for Networks. In proceedings of LISA '99: 13th Systems Administration Conference (pp. 229-238). Seattle, WA : USENIX.
- Sidhu, R. & Prasanna, V.K. (2001). Fast Regular Expression Matching using FPGAs. In proceedings of the 9th International IEEE symposium on FPGAs for Custom Computing Machines, FCCM'01. Rohnert Park, California, USA.
- Sugawara, Y., Inaba, M., & Hiraki, K. (2004). Over 10 Gbps String Matching Mechanism for Multi-stream Packet Scanning Systems. In proceedings of Field Programmable Logic and Applications, 14th International Conference, FPL 2004 (pp. 484-493). Springer-Verlag.
- Tan, L., & Sherwood, T. (2005). A High Throughput String Matching Architecture for Intrusion Detection and Prevention. In the proceedings of the 32nd Annual International Symposium on Computer Architecture (ISCA 2005). Madison, Wisconsin, USA.
- Tripp, G. (2005). A finite-state-machine based string matching system for intrusion detection on high-speed networks. In Paul Turner and Vlasti Broucek, editors, EICAR Conference Best Paper Proceedings, (pp. 26-40). Saint Julians, Malta : EICAR.
- Xilinx Virtex-II Platform FPGAs: Complete Data Sheet – Product Specification. (2005). Xilinx Inc. Retrieved 9 March 2006 from <http://direct.xilinx.com/bvdocs/publications/ds031.pdf>

Author's Errata

In the first row of data in Table 5, the value of **12** for the byte position where a match of the string “cybercop” occurs is incorrect – this value should be **11**.

This error does NOT appear in the extended version of this paper that was published in the special issue of “Journal in Computer Virology”.