

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 March 2003 (20.03.2003)

PCT

(10) International Publication Number
WO 03/023548 A2

- (51) International Patent Classification⁷: **G06F**
- (21) International Application Number: PCT/US01/45817
- (22) International Filing Date: 3 December 2001 (03.12.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/322,012 12 September 2001 (12.09.2001) US
- (71) Applicant: **RAQIA NETWORKS, INC.** [US/US]; 181 Wells Ave., Newton, MA 02459 (US).
- (72) Inventors: **WYSCHOGROD, Daniel**; 28 Judith Road, Newton, MA 02459 (US). **ARNAUD, Alain**; 24 Cummings Road, Newton, MA 02459 (US). **LEES, David, Eric, Berman**; 57 Gleason Road, Lexington, MA 02420 (US). **LEIBMAN, Leonid**; 40 Kendal Common Rd., Weston, MA 02493 (US).
- (74) Agent: **GALBI, Elmer**; 13314 Vermeer Drive, Lake Oswego, OR 97035 (US).


- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: HIGH SPEED DATA STREAM PATTERN RECOGNITION



(57) Abstract: A system and method in accordance with the present invention determines in real-time the portions of a set of characters from a data or character stream which satisfies one or more predetermined regular expressions. A Real-time Deterministic Finite state Automaton (RDFA) ensures that the set of characters is processed at high speeds

with relatively small memory requirements. An optimized state machine models the regular expression(s) and state related alphabet lookup and next state tables are generated. Characters from the data stream are processed in parallel using the alphabet lookup and next state tables, to determine whether to transition to a next state or a terminal state, until the regular expression is satisfied or processing is terminated. Additional means may be implemented to determine a next action from satisfaction of the regular expression.



WO 03/023548 A2

HIGH SPEED DATA STREAM PATTERN RECOGNITION

Field Of The Invention:

The present invention generally relates to systems and methods for performing, at high speeds, pattern recognition from streams of digital data.

Background Of The Invention:

With the continued proliferation of networked and distributed computers systems, and applications that run on those systems, comes an ever increasing flow and variety of message traffic between and among computer devices. As an example, the Internet and world wide web (the "Web") provide a global open access means for exchanging message traffic. Networked and/or distributed systems are comprised of a wide variety of communication links, network and application servers, sub-networks, and internetworking elements, such as repeaters, switches, bridges, routers, gateways.

Communications between and among devices occurs in accordance with defined communication protocols understood by the communicating devices. Such protocols may be proprietary or non-proprietary. Examples of non-proprietary protocols include X.25 for packet switched data networks (PSDNs), TCP/IP for the Internet, a manufacturing automation protocol (MAP), and a technical & office protocol (TOP). Other proprietary protocols may be defined as well. For the most part, messages are comprised of packets, containing a certain number of bytes of information. The most common example is Internet Protocol (IP) packets, used among various Web and Internet enabled devices.

1

2 A primary function of many network servers and other network devices (or nodes),
3 such as switches, gateways, routers, load balancers and so on, is to direct or
4 process messages as a function of content within the messages' packets. In a
5 simple, rigid form, a receiving node (e.g., a switch) knows exactly where in the
6 message (or its packets) to find a predetermined type of contents (e.g., IP
7 address), as a function of the protocol used. Typically, hardware such as switches
8 and routers are only able to perform their functions based on fixed position
9 headers, such as TCP or IP headers. Further, no deep packet examination is
10 done. Software, not capable of operating at wire speed is sometimes used for
11 packet payload examination. This software does not typically allow great flexibility
12 in specification of pattern matching and operates at speeds orders of magnitude
13 slower than wire rate. It is highly desirable to allow examination and recognition
14 of patterns both in packet header and payload described by regular expressions.
15 For example, such packet content may include address information or file type
16 information, either of which may be useful in determining how to direct or process
17 the message and/or its contents. The content may be described by a "regular
18 expression", i.e., a sequence of characters that often conform to certain
19 expression paradigms. As used herein, the term "regular expression" is to be
20 interpreted broadly, as is known in the art, and is not limited to any particular
21 language or operating system. Regular expressions may be better understood
22 with reference to Mastering Regular Expressions, J. E. F. Friedl, O'Reilly,
23 Cambridge, 1997.

24

25 It is clear that the ability to match regular expressions would be useful for content

1 based routing. For this, a deterministic finite state automaton (DFA) or non-
2 deterministic finite state automaton (NFA) would be used. The approach used
3 here follows a DFA approach. A conventional DFA requires creation of a state
4 machine prior to its use on a data (or character) stream. Generally, the DFA
5 processes an input character stream sequentially and makes a state transition
6 based on the current character and current state. This is a brute-force, single
7 byte at a time, conventional approach. By definition, a DFA transition to a next
8 state is unique, based on current state and input character. For example, in prior
9 art FIG. 1A, a DFA state machine 100 is shown that implements a regular
10 expression "binky.*\jpg". DFA state machine 100 includes states 0 through 9,
11 wherein the occurrence of the characters 110 of the regular expression effect the
12 iterative transition from state to state through DFA state machine 100. The start
13 state of the DFA state machine is denoted by the double line circle having the
14 state number "0". An 'accepting' state indicating a successful match is denoted
15 by the double line circle having the state number "9". As an example, to transition
16 from state 0 to state 1, the character "b" must be found in the character stream.
17 Given "b", to transition from state 1 to state 2, the next character must be "i".
18
19 Not shown explicitly in FIG. 1A are transitions when the input character does not
20 match the character needed to transition to the next state. For example, if the
21 DFA gets to state 1 and the next character is an "x", then failure has occurred and
22 transition to a failure terminal state occurs. FIG. 1B shows part 150 of FIG. 1A
23 drawn with failure state transitions, wherein a failure state indicated by the "Fail"
24 state. In FIG. 1B, the tilde indicates "not". For example, the symbol "~b" means
25 the current character is "not b". Once in the failure state, all characters cause a

1 transition which returns to the failure state, in this case.

2

3 Once in the accepting state, i.e., the character stream is "binky.*\jpg", the
4 receiver node takes the next predetermined action. In this example, where the
5 character stream indicates a certain file type (e.g., ".jpg"), the next predetermined
6 action may be to send the corresponding file to a certain server, processor or
7 system.

8

9 While such DFAs are useful, they are limited with respect to speed. The speed of
10 a conventional DFA is limited by the cycle time of memory used in its
11 implementation. For example, a device capable of processing the data stream
12 from an OC-192 source must handle 10 billion bits/second (i.e., 10 gigabits per
13 second (Gbps)). This speed implies a byte must be processed every 0.8
14 nanosecond (nS), which exceeds the limit of state of the art memory. For
15 comparison, high speed SDRAM chips implementing a conventional DFA operate
16 with a 7.5 nS cycle time, which is ten times slower than required for OC-192. In
17 addition, more than a single memory reference is typically needed, making these
18 estimate optimistic. As a result, messages or packets must be queued for
19 processing, causing unavoidable delays.

20

21 **Summary Of The Invention:**

22 A system and method in accordance with the present invention determines in
23 real-time whether a set of characters from a data or character stream (collectively
24 "data stream") satisfies one or more of a set of predetermined regular
25 expressions. A regular expression may be written in any of a variety of codes or

1 languages known in the art, e.g., Perl, Python, Tcl, grep, awk, sed, egrep or
2 POSIX expressions. Additional means may be implemented to determine a next
3 action from satisfaction of one such regular expression, or from a lack of such
4 satisfaction of a regular expression. The present invention provides, an improved
5 high speed, real-time DFA, called a **R**Real-time **D**Deterministic **E**inite state
6 **A**utomaton (hereinafter RDFA). The RDFA provides high speed parallel pattern
7 recognition with relatively low memory storage requirements. The RDFA includes
8 a DFA optimized in accordance with known techniques and a set of alphabet
9 lookup and state related tables that combine to speed up the processing of
10 incoming data streams.

11

12 The data stream may be received by a typical computer and/or network device,
13 such as a personal computer, personal digital assistant (PDA), workstation,
14 telephone, cellular telephone, wireless e-mail device, pager, network enabled
15 appliance, server, hub, router, bridge, gateway, controller, switches, server load-
16 balancers, security devices, nodes, processors or the like. The data stream may
17 be received over any of a variety of one or more networks, such as the Internet,,
18 intranet, extranet, local area network (LAN), wide area network (WAN), telephone
19 network, cellular telephone network, and virtual private network (VPN).

20

21 An RDFA system in accordance with the present invention includes a RDFA
22 compiler subsystem and a RDFA evaluator subsystem. The RDFA compiler
23 generates a set of tables which are used by the RDFA evaluator to perform
24 regular expression matching on an incoming data stream. The data stream may
25 present characters in serial or parallel. For example, four characters at a time

1 may arrive simultaneously or the four characters may be streamed into a register.

2 The RDFA evaluator is capable of regular expression matching at high speed on
3 these characters presented in parallel.

4
5 The RDFA compiler subsystem generates a DFA state machine from a user
6 specified regular expression. The DFA state machine is optimized to include a
7 minimum number of states, in accordance with known techniques. A number of
8 bytes to be processed in parallel is defined, as M bytes. For each state in the
9 state machine, the RDFA compiler determines those characters, represented by
10 bytes, that cause the same transitions. Those characters that cause the same
11 transitions are grouped into a class. Therefore, each class, for a given current
12 state of the state machine, includes a set of characters that all cause the same
13 transitions to the same set of next states. Each class is represented by a class
14 code. The number of bits required for a class code is determined solely from the
15 number of classes at a given state and byte position.

16
17 The RDFA compiler generates a set of state dependent alphabet lookup tables
18 from the class codes and the relevant alphabet. A lookup table for a given state
19 associates a class code to each character in the relevant alphabet. Each of M
20 bytes under evaluation has its own lookup table. Classes are a compressed
21 representation of the alphabet used in a state machine, since multiple symbols
22 can be represented by a single class. This can lead to large reductions in the
23 number of bits required to represent alphabet symbols, which in turn leads to
24 large reductions in the size of next state lookup tables.

25

1 The RDFFA compiler then generates a set of state dependent next state tables
2 using the class codes and state machine. For each state (as a current state) in
3 the state machine, a set of next states is determined and represented in a next
4 state table. For a given current state, the next state is a function of the characters
5 represented by the bytes under evaluation. The possible sets of concatenated
6 class codes from the alphabet lookup tables serve as indices to the possible next
7 states in the appropriate next state table. Given a current state, a table of
8 pointers may be defined, wherein each pointer points to the appropriate next state
9 table from the set of next state tables. The RDFFA compiler can also determine
10 the memory requirements for RDFFA system data associated with a defined
11 regular expression.

12

13 During parallel evaluation, the RDFFA evaluator selects, or accepts, the next M
14 bytes and gets the appropriate M lookup tables to be applied to the bytes under
15 evaluation. Each byte is looked up in its corresponding lookup table to determine
16 its class code. As previously mentioned, the class codes are concatenated.
17 Given a current state, the RDFFA evaluator retrieves the appropriate next state
18 table. The code resulting from concatenation of the class code lookup results, is
19 applied as an index to the selected next state table to determine the next state
20 which involves M transitions beyond the current state.

21

22 This process continues until evaluation is terminated or the regular expression is
23 satisfied. The process may be terminated when, for example, the bytes under
24 evaluation do not cause a transition to a non-failure state. With a regular
25 expression satisfied, the next action may be determined by the RDFFA system, or

1 by a system interfaced therewith. The RDFFA system may be employed in any of
2 a variety of contexts where it is essential or desirable to determine satisfaction of
3 a regular expression, whether anchored or unanchored, in a data stream,
4 particularly when such determinations are to be made at high speeds, such as
5 required by OC-192 rates. The RDFFA system may also be employed in contexts
6 where consumption of relatively small amounts of memory by the RDFFA system
7 data are required or desirable.

8 9 **Brief Description Of The Drawings**

10 The foregoing and other objects of this invention, the various features thereof, as
11 well as the invention itself, may be more fully understood from the following
12 description, when read together with the accompanying drawings, described:

13
14 FIG. 1A is a state diagram implementing a regular expression, in accordance with
15 the prior art;

16 FIG. 1B is a portion of the state diagram of the regular expression of FIG. 1A,
17 including a failure state;

18 FIG. 2A is a block diagram of a RDFFA system in accordance with the present
19 invention;

20 FIG. 2B is a block diagram of a RDFFA compiler, from the RDFFA system of FIG.
21 2A;

22 FIG. 2C is a block diagram of a RDFFA evaluator, from the RDFFA system of FIG.
23 2A;

24 FIG. 3 is a diagram depicting 4 byte parallel processing and 4 corresponding
25 alphabet lookup tables, used by the RDFFA evaluator of FIG. 2C;

1 FIG. 4 is a diagram depicting a next state table, used by the RDFA evaluator of
2 FIG. 2C;

3 FIG 4A is a diagram indicating the flow of data from the character tables, the
4 index table and memory.

5 FIG. 5 is a diagram depicting characters that cause the same state transitions,
6 used by the RDFA compiler of FIG. 2B; and

7 FIG. 6 is a diagram depicting a state machine used by the RDFA compiler of FIG.
8 2B.

9 FIG. 6A illustrates a number of states reachable by 2-closure.

10 FIG 7A illustrates a DFA for processing 8 characters.

11 FIG 7B illustrates an RDFA for processing 4 bytes in parallel

12
13 **Detailed Description Of The Preferred Embodiments:**

14 In the preferred embodiment, the present invention is implemented as a RDFA
15 system 200, shown in FIG. 2A, which includes two subsystems. The first
16 subsystem is a RDFA compiler 210 that performs the basic computations
17 necessary to create tables for subsequent real-time pattern recognition. The
18 second subsystem is a RDFA evaluator 250 that performs the evaluation of
19 characters using the RDFA tables created by the RDFA compiler 210. The RDFA
20 system 200 includes a first memory 220 for high speed access by RDFA
21 evaluator 250 during evaluation of characters from the data stream. This first
22 memory 220 consists of on-chip or off-chip or any combination thereof. A second
23 memory 204 includes the initial one or more regular expressions of interest, and
24 need not lend itself to high speed access, unless required as a function of a
25 particular application to which the RDFA is applied.

1

2 FIG. 2B is a block diagram of the RDFFA compiler 210. As will be discussed in
3 more detail below, the RDFFA compiler 210 includes a regular expression compiler
4 212 that converts a regular expression, from memory 204, into an optimized state
5 machine. An alphabet lookup table generator 214 generates, from the regular
6 expression and the state machine, a series of state dependent alphabet lookup
7 tables. The alphabet lookup tables include codes associated with each character
8 in an applicable alphabet of characters. These alphabet lookup tables are stored
9 in high speed memory 220. During RDFFA data stream processing (i.e., character
10 evaluation), a character represented by a byte under evaluation is looked up in a
11 corresponding alphabet lookup table to determine its state dependent code, as
12 will be discussed in greater detail. A next state table generator 216 generates a
13 table of next states of the state machine to be applied during evaluation of a set
14 of characters, wherein next states are determined as a function of a current state
15 and the character codes from the alphabet lookup tables. The next state table is
16 also preferably stored in high speed memory 220.

17

18 FIG. 2C is a functional block diagram of the RDFFA evaluator 250. The RDFFA
19 evaluator 250 includes several functional modules that utilize the alphabet lookup
20 tables and next state tables generated by the RDFFA compiler 210. At a top level,
21 a byte selector module 252 captures the requisite number of bytes (i.e., M bytes)
22 from an incoming data stream 205. An optional bit mask 251 can filter the input
23 stream to select words from predetermined positions, allowing the processing to
24 ignore certain portions of the input stream. Each bit in the mask corresponds to a
25 four byte section of a packet for this embodiment. The selected bytes are taken

1 and processed in parallel by an alphabet lookup module 254, which selectively
2 applies the alphabet lookup tables from memory 220 to determine a character
3 class code for each byte. As will be discussed in greater detail, characters
4 causing the same state transition are grouped in classes, which are represented
5 in alphabet lookup tables as class codes of a certain bit width. The alphabet
6 lookup module 254 concatenates the class codes obtained from the lookup tables
7 and passes the concatenated code to a next state module 256. The next state
8 module 256 selectively applies the concatenated class codes to the appropriate
9 next state table from memory 220, given a current state, to determine a next *M*th
10 state in a corresponding state machine. This process continues at least until a
11 failure state or accepting state is achieved.

12
13 The RDFA evaluator 250, as well as the RDFA compiler 210, may be
14 implemented in hardware, software, firmware or some combination thereof. In the
15 preferred form, the RDFA evaluator 250 is a chip-based solution, wherein RDFA
16 compiler 210 and high speed memory 220 may be implemented on chip 270.
17 Memory 204 may also be on-chip memory or it may be off-chip memory, since
18 high-speed is typically not as vital when generating the RDFA. However, if high-
19 speed is required the RDFA compiler 210 and memory 204 may each be on-chip.
20 Therefore, preferably, to achieve higher speeds the primary functionality of RDFA
21 evaluator 250 for processing incoming data streams is embodied in hardware.
22 The use of pointers to next state tables, rather than directly using the alphabet
23 table lookup results, allows flexibility in memory management. For example, if on-
24 chip and off-chip memory is available, then pointers can be used so that more
25 frequently used memory is on-chip, to speed up RDFA performance. The RDFA

1 expression compiler 210 will determine the amount of memory required. This
2 allows the user to know if a particular set of rules will fit in the on-chip memory.
3 Thus, memory related performance can be accurately known ahead of time.

4
5 As will be appreciated by those skilled in the art and discussed in further detail
6 below, a RDFFA system 200 in accordance with the present invention requires
7 relatively modest amounts of high speed or on-chip memory 220, certainly within
8 the bounds of that which is readily available. Memory 220 is used to store the
9 alphabet lookup tables and next state tables for a given regular expression.

10 Unlike a conventional (i.e., single byte at a time processing) DFA approach, a
11 RDFFA is configured for scalable parallel processing. As a general rule, increasing
12 the number of bytes (M) processed in parallel yields increasingly greater
13 processing speeds, subject to the limitations of other relevant devices. While in
14 the preferred embodiment provided herein, the RDFFA evaluator 250 processes
15 four (4) bytes in parallel (i.e., $M = 4$), there is no inherent limitation to the number
16 of bytes that can be processed in parallel.

17
18 **Data Stream Evaluation:**

19 FIG. 3 illustrates the multiple alphabet lookup table concept 300 for a set of 4
20 bytes 320, which are selected from the data stream 205 by byte selector 252 and
21 are taken as parallel input 260 by alphabet lookup module 254 (see FIG. 2C).

22 Each byte represents a character (e.g., a number, a letter, or a symbol) from the
23 permitted alphabet. In the preferred embodiment, a separate alphabet lookup
24 table having 256 elements is defined for each of the 4 bytes and each state and is
25 stored in memory 220. The alphabet lookup tables 310 are formed and applied

1 as a function of a current state of a state machine that represents the regular
2 expression.

3
4 In the example of FIG. 3, a first alphabet lookup table 312, having a 2 bit width, is
5 used to lookup a first byte 322. A second alphabet lookup table 314, having a 3
6 bit width, is used to lookup a second byte 324, and so forth with alphabet tables
7 316 and 318 and third byte 326 and fourth byte 328, respectively. The elements
8 of the alphabet lookup tables 310 are related to state transitions for a
9 corresponding state machine that models the regular expression. Accordingly,
10 the selection and application of alphabet lookup tables 310 is a function of the
11 current state of the state machine. The current state is the last state resulting
12 from the processing of the previous 4 characters, if any. Thus, a different set of
13 alphabet lookup tables is used for each current state.

14
15 The widths of the table entries for each byte can vary from one state to the next,
16 depending on the regular expression and the current state of the corresponding
17 state machine. In the FIG. 3 example, the table widths in bits are 2 bits for table
18 312, 3 bits for table 314, 3 bits for table 316, and 4 bits for table 318. The table
19 widths in another state might be 1 bit, 1 bit, 2 bits, and 4 bits, as an example. For
20 instance, if for the first byte there are only two possible character classes, then
21 the width of the alphabet lookup table for that bit need only be 1 bit. The current
22 state is stored in memory (e.g., on-chip memory 220) for use in determining which
23 alphabet lookup tables to apply to the 4 bytes 320 and for determining a next
24 state.

25

1 For each of the 4 bytes 320, using lookup tables 310 a different class code is
2 obtained by alphabet lookup module 254. As previously discussed, the
3 characters are grouped into classes according to the state transitions the
4 characters cause and codes associated with those classes (i.e., class codes) are
5 represented in the alphabet lookup tables. Therefore, if byte 322 represents the
6 character "a", alphabet lookup module 254 finds the element in alphabet lookup
7 table 312 that corresponds to "a" and obtains the class code stored at that
8 element (e.g., class code 01). This is done for each other byte (i.e., bytes 324,
9 326 and 328) using their respective alphabet lookup tables (i.e., tables 314, 316
10 and 318).

11

12 The lookup table class codes for each of the 4 bytes are concatenated together,
13 which for the FIG. 3 example produces a 12 bit result (i.e., 2 + 3 + 3 + 4 bits). As
14 an example, assume that from lookup tables 310 of FIG. 3 resulted a 2 bit word
15 "01" from table 312, a 3 bit word "001" from table 314, a 3 bit word "011" from
16 table 316, and a 4 bit word "0000" from table 318. The resulting 12 bit
17 concatenated word would be "010010110000".

18

19 As is shown in FIG. 4, the current state of the state machine is used as an index
20 into a table of pointers 410. Table 410 is defined as a function of the regular
21 expression's state machine, so each current state has a corresponding table to
22 possible next states. Each pointer in table 410 points to a linear (i.e., 1
23 dimensional (1-D)) table 420 of next state values (or a "next state table") and the
24 12 bit concatenated result of the parallel alphabet lookup is used as an offset or
25 index into the selected next table 420. Therefore, a next state value is selected

1 from next state table 420 as a function of the current state and the concatenated
2 12 bit word. The selected next state value corresponds to the next state. The
3 next state determined from evaluation of the 4 bytes serves as the current state
4 for evaluation of the next set of 4 bytes.

5

6 In the preferred form, the selected next state table value includes a terminal state
7 code (e.g., with higher order bit set to 1) that indicates whether or not the next
8 state is an accepting state (or terminal state). Generally, a terminal state is a
9 state the process enters when processing from a data stream with respect to a
10 certain one or more regular expressions is completed; i.e., it is indicative of
11 termination of processing with respect to the one or more regular expressions.

12 For example, in the preferred embodiment a high order bit associated with one or
13 more of the bytes under evaluation is set to "1" upon transition into a terminal
14 state. In one embodiment, the hardware stores the word (i.e., the 4 bytes under
15 evaluation) for which the terminal state occurred and the corresponding offset
16 from the lookup table (i.e., the 12 bit concatenated word). Thereafter, post
17 processing software may use the stored data to determine at which of the 4 bytes
18 the regular expression terminated. This is useful in many situations where only a
19 small number of regular expression matches occur per packet, so the number of
20 such determinations is relatively small. In another embodiment, the codes (i.e.,
21 the 4 bytes and 12 bit word) are stored in a secondary terminal state table, which
22 allows the hardware to directly determine which byte terminated the processing.
23 The benefit of allowing the hardware to make such determinations is that it can be
24 accomplished much more quickly in hardware, which is a significant consideration
25 in high speed, real-time processing.

1

2 In accordance with the preferred embodiment, only three (3) memory operations
3 are required to process the 4 bytes. They are: (i) find characters in lookup tables
4 310; (ii) find pointer in table 410; and (iii) get next state indicia from next state
5 table 420. Further, these operations may be easily pipelined by performing the
6 character table lookup at the same time as the last 4 byte result is being looked
7 up in the next state table. to allow improved processing times , with the only
8 significant limitation being the longest memory access.

9

10 The benefits of the preferred embodiment can be further appreciated when the
11 RDFA memory requirements are compared with those of a naïve DFA approach,
12 where the lookup is applied to a 4 byte word. In this type of DFA parallelization, 4
13 bytes would be looked up in parallel. This would require a table having 256^4
14 entries, which is about 4.295 billion entries, and a word (4 byte) cycle time of 3.2
15 nS in order to keep up with OC-192 rates (i.e., 10 Gb/sec). This is impractical to
16 implement with current or near-term memory technology, based on the speed and
17 size required to keep up with OC-192 rates. Further, such a large amount of
18 memory cannot be implemented on-chip, so a significant amount of off-chip
19 memory would be required, unacceptably slowing the process. Compare the
20 memory requirement of simple DFA parallelization with the greatly reduced
21 amount of memory used in the preferred embodiment of the RDFA system 200.
22 Note that the naïve DFA parallelization requires many orders of magnitude
23 greater memory size than an RDFA system 200, in accordance with the present
24 invention.

25

1 Fig. 4A is another illustration of the alphabet lookup tables 310, the index table
2 410, and the next state table 450 showing their operation and interactions. The
3 bytes which are being examined are designated 320. Bytes 320 are a four byte
4 segment from a data stream (i.e. four bytes from a data packet that is being
5 examined). The alphabet lookup tables 310 have a segment associated with
6 each possible state of the state machine. In Fig. 4A the states are designated s1,
7 s2, s3, etc. along the left side of the figure. In each state, the bytes 320 are used
8 to interrogate the section of table 310 associated with that particular state. The
9 lookup operation produces a multi-bit result 330. The number of bits in the result
10 330 (i.e. the number of bits retrieved or generated by the alphabet lookup table) is
11 a function of the particular bytes 320, the particular state, and the byte position.
12 The index table 410 has an entry for each state. Each entry in table 410 includes
13 a code which tells the system how many bits to retrieve from the output of the
14 lookup table 310 for that particular state. (Note in an alternate embodiment this
15 code is stored in a separate table that has a location for each state similar to table
16 410) During any particular state conventional addressing circuits address and
17 read the contents of the location in table 410 associated with the particular state.
18 The result bits 330 and the bits in the current state position of the index table 410
19 are concatenated to produce a memory address 441. As indicated at the left side
20 of Fig. 4A, different locations in index table 410 have a different number of bits.
21 The number of bits in concatenated result 441 (i.e. total number of bits from table
22 410 and result 330) is always a fixed number. In the preferred embodiment that
23 number is 13. The number of bits equals the number of bits in a memory address
24 for the particular system. Thus, for each state the associated location in table
25 410 indicates how many bits should be retrieved from table 310 and it provides a

1 series of bits so that there is a total of 13 bits for address 441.

2

3 Address 441 is the address of an entry in the next state table 450. The memory
4 address 441 is used to interrogate next state table 450 utilizing conventional
5 memory addressing circuitry. The entry in next state table 450 at address 441
6 indicates the next state. The entry in the next state table 450 may also contain a
7 flag which indicates that the operation has reached a special point such as a
8 termination point.

9

10 The operations proceed until the flag in the next state table indicates that the
11 operation has reached a termination point or that the bytes have been recognized
12 or matched. When a match is found, processing the bytes in a particular packet
13 can then either terminate or the system can be programmed to continue
14 processing other sets of bytes 320 in an attempt to find other matching patterns.

15

16 If the next state table does not indicate that the operation has terminated, the
17 process proceeds to the next state and the process repeats. If the process
18 repeats the information in appropriate next state table 450 is used. That is, the
19 designation of the next state in table 450 is used to generate the address of an
20 appropriate section of lookup table 310 and the process repeats. Upon reaching
21 a termination state, the following data is saved in memory registers 442:

22

1. Pointer to the word (4 bytes) in the packet at which the terminal state
23 occurred.

23

24 2. The table offset (computed from the alphabet table lookups results and
25 index table) into the next-state table.

25

1 The saved data can be used by post processing operations which determine what
2 action to take after the operation has terminated. In some embodiments when a
3 termination flag is encountered which indicates that a match is found, the
4 operation continues, that is, additional bytes in the string is are processed in an
5 effort to locate another match to the specified regular expression.

6
7 In general after four bytes have been processed, four different bytes are
8 streamed into register 320 and the process repeats. Furthermore, one can
9 search for a wide array of different patterns. A target pattern can be more than
10 four bytes long. For example if one is searching for a five byte pattern, after four
11 of the bytes have been located another set of four bytes can be streamed into
12 register 320 to see if the fifth byte is at an appropriate location.

13
14 It is noted that each different set of regular expressions which one wants to locate
15 require a different set of data in tables 310, 410 and 450. It should be
16 understood that the invention can be used to locate any desired regular
17 expression, not just fixed character sequences.

18
19 As one particular example wherein the tables are provided with data to recognize
20 the particular character string "raqia " the tables provide for 32 states of operation.

21 The four tables 310 each have 32 sections each with 256 entries for a total of
22 8192 entries. The index table has 32 entries. It is noted that the choice of 32
23 states is matter of engineering choice for the particular example. In the particular
24 example given above the next state table 450 has 8192 entries. It is noted that
25 the number of entries in this table is also a matter of choice. The number of

1 entries in the next state table for each state is determined by the number of
2 combinations of character classes for that state for all the byte positions. For
3 example, if the number of character classes for byte positions 0 through 3 are 4,
4 4, 8, 8 respectively, then the total number of next state table entries for that state
5 is $4 \times 4 \times 8 \times 8 = 1024$. And the total size of the address space for all the states is
6 the sum of the table sizes for each state. In one embodiment the number of
7 character classes at each byte position is a power of 2, but other embodiments
8 use various different numbers of character classes.

9

10 It should be noted that while in the embodiment described, four bytes are
11 processed in parallel, alternate embodiments can be designed to handle different
12 numbers of bits in parallel. For example other embodiments can handle 1, 2, 6,
13 8, 12 bytes in parallel.

14

15 Creation of the RDFA Tables:

16 To generate a RDFA in accordance with the present invention, the regular
17 expression compiler 212 converts a regular expression from memory 204 into a
18 DFA. The regular expression compiler 212 may also optimize the DFA to
19 minimize the number of states. These processes are known in the art, so are not
20 discussed in detail herein. The regular expression compiler is also configured to
21 determine the amount of memory required to store the RDFA for a given regular
22 expression, as will be discussed in further detail below. This allows the user to
23 know if a particular set of rules (i.e., regular expressions) will fit in the on-chip
24 memory. Thus, performance can be accurately predicted.

25

1 The regular expression compiler 212 also reduces state transition redundancy in
2 the alphabet representing the input data stream by recognizing that DFA state to
3 state transition decisions can be simplified by grouping characters of an alphabet
4 according to the transitions they cause. The list of states that may be reached in
5 a single transition is referred to as '1-closure'. The term " n -closure" is defined as
6 the list of states reachable in n transitions from the current state. n -closure is
7 readily calculated recursively as the list of states reachable from the $n-1$ closure.
8 There may be more than one character that causes the same transitions to the
9 same n -closure set. In such a case, characters may be grouped into classes
10 according to the set of next state transitions they cause. Rather than
11 representing individual characters, each class may be represented in a 1, 2, 3, or
12 4 bit code, for example. In this manner, the applicable alphabet is represented in
13 an extremely economical form.

14
15 Even very complicated expressions can achieve significant compression in the
16 number of bits required to represent its alphabet by mapping to character classes.
17 For example, a portion of a regular expression represented as "(a | b | c | g)" can
18 be represented in a state transition diagram 500, shown in FIG. 5, wherein the
19 expression indicates "a" or "b" or "c" or "g". These characters all cause a
20 transition from state "1" to state "2", therefore, these characters can all be
21 mapped into a single class. If all other characters cause a transition to a failure
22 state, then all of those characters can be grouped into a second class. Therefore,
23 when in state 1 of FIG. 5 (i.e., state 1 is the current state) all transitions can be
24 represented with a 1 bit code, wherein a code of "1" could indicate a transition
25 into state "2" and a code of "0" could indicate a transition into a failure state F (not

1 shown). Mapping characters into classes eliminates the need to represent
2 characters with 8 bits, as is typical with conventional approaches.

3

4 Alphabet lookup tables are generated by the alphabet table generator 214 of
5 FIG. 2B. In the present invention, the RDFA alphabet lookup tables reflect the
6 classes that represent the state transitions, wherein the 1, 2, 3, or 4 bit
7 representation, as the case may be, are embodied in character classes related to
8 the current state. In general, when M bytes are processed in parallel a separate
9 alphabet lookup table is computed for each of the M bytes. Further, a different
10 set of tables is computed for each state in the state machine. Thus, if a state
11 machine has L states and M bytes are processed in parallel, a total of $(L \times M)$
12 alphabet lookup tables are produced, in accordance with the preferred
13 embodiment.

14

15 The algorithm used to produce the M character class tables for a regular
16 expression state machine from a starting state S , is as follows. The n th alphabet
17 lookup table (where $1 \leq n \leq M$) uses the previously computed $n-1$ closure and
18 then computes the n -closure. Then, for each character in the alphabet, a list of
19 non-failure state transitions from the $n-1$ closure to the n -closure is generated. An
20 alphabet mapping is then initialized by placing the first character in the alphabet
21 into character class 0. The transition list for the next letter in the alphabet, for a
22 given regular expression, is examined and compared with the transitions for the
23 character class 0. If they are identical, then the character is mapped to class 0,
24 otherwise a new class called "class 1" is created and the character is mapped to
25 it. This process proceeds for each character in the alphabet. So, if a list of

1 transitions for a character matches the transitions for an existing class, then that
2 character is represented in that existing class, otherwise that character is the first
3 member of a new class. The result of this process is a character class number
4 associated with each character in the alphabet. The total number of classes for a
5 particular lookup table may be represented by P . Then, the number of bits
6 necessary to represent each symbol is given by:

$$7 \quad Q = \text{floor}(\log_2 P) + 1$$

8 Q is also the width of the table entries in the alphabet lookup table (e.g., 1, 2, 3, or
9 4 bits). For example, in alphabet lookup table 312 of FIG. 3, $Q = 2$. Note that Q
10 is computed for each alphabet lookup table separately and varies as a function of
11 both state and byte position.

12

13 This concept may be appreciated with a simple example for processing 2 bytes in
14 parallel (i.e., for $M = 2$) for the portion 650 of a state machine 600 shown in FIG.
15 6. This example focuses primarily on lookup tables and transitions from the 0
16 state. State machine 600 is derived from a predefined regular expression. The 1-
17 closure for state 0 is (1, 2, 3, F), where the failure (or terminal) state may be
18 denoted by symbol F (not shown). That is, as can be seen from FIG. 6, from
19 state 0, non-failure transitions may be to state 1, state 2, or state 3. Table 1
20 provides a list of state transitions out of state 0 for an alphabet consisting of
21 letters from "a" to "k", in accordance with the state diagram 600 of FIG. 6.

22

23

24

Letter	Transitions
a	$0 \Rightarrow 1$
b	$0 \Rightarrow 1$
c	$0 \Rightarrow 2$
d	$0 \Rightarrow 2$
e	$0 \Rightarrow 3$
f	$0 \Rightarrow 3$
g	$0 \Rightarrow 3$
h	\emptyset null (or <i>F</i> state)
i	\emptyset null (or <i>F</i> state)
j	\emptyset null (or <i>F</i> state)
k	\emptyset null (or <i>F</i> state)

1 **Table 1 - Transitions for State Diagram 600, From State 0**

2 Upon inspection, Table 1 shows that the alphabet maps to 4 different equivalent
3 classes, meaning that 2 bits are sufficient for the width of an alphabet lookup
4 table for a current state of state 0. Therefore, with regard to a current state 0, the
5 following classes may be formed: class 0 (a, b), class 1 (c, d), class 2 (e, f, g,) and
6 a failure state class 3 (h, i, j, k). In the corresponding alphabet lookup table, class
7 0 may be represented as "00", class 1 as "01", class 2 as "10", and class 3 as
8 "11", as follows:

9
10
11
12

Letter	Class Code In Alphabet Lookup Table
a	00
b	00
c	01
d	01
e	10
f	10
g	10
h	11
l	11
j	11
k	11

1 **Table 2 - Lookup Table Entries, Current State 0**

2 The 2-closure for state machine 600 is (1, 4, 5, 6, 7, 8, F) from state 0. Similarly,
3 Table 3 is a list of state transitions for each character for the 2-closure. In this
4 case, inspection of Table 3 shows the alphabet maps to 8 equivalent characters,
5 so that 3 bits are required for the table width. Note that as indicated in the Q
6 value calculation, if the number of equivalent characters had been 5, the table
7 width would still be 3 bits.

8

9

10

11

12

Letter	Transitions
a	1 \Rightarrow 4
b	\emptyset null (or <i>F</i> state)
c	1 \Rightarrow 1
d	1 \Rightarrow 4 2 \Rightarrow 5
e	3 \Rightarrow 6
f	\emptyset null (or <i>F</i> state)
g	\emptyset null (or <i>F</i> state)
h	2 \Rightarrow 6
l	3 \Rightarrow 7
j	3 \Rightarrow 8
k	\emptyset null (or <i>F</i> state)

Table 3 - Transitions for State Diagram 600, From State 1

1
2 The next state table generator 216 of FIG. 2B generates, for each state in the
3 state machine (as a current state), a list of next states represented in next state
4 tables (e.g., next state table 420 of FIG. 4). Therefore, for each current state,
5 there is a one-dimensional (1-D) list of possible next states. Using state machine
6 600 of FIG. 6, assume the current state is state 0 and $M = 4$ (where M is the
7 number of bytes processed in parallel). As was discussed with respect to states 0
8 through 3 above, and shown in Tables 1 - 3, the characters that cause state
9 transitions can be mapped into classes. The corresponding next state table is
10 comprised of entries that dictate the next state given the class codes of the four
11 characters (or bytes) under evaluation from the alphabet lookup tables.

12

1 Assume 4 bytes were received representing the 4 characters "c, h, i, e". As
2 mentioned previously, the class code for "c" with a current state 0 is 01. The
3 class code for "h" is 3 bits, as the second of 4 bytes assume its class code is 011.
4 Also, assume for a current state 0 and as the third of 4 bytes, the class code for
5 "i" is 0011. Finally, assume for a current state of 0 and as the fourth byte, the
6 class code for "e" is 101. The corresponding next state table, will have a next
7 state value corresponding to state 12, given the above class codes for the 4 bytes
8 and a current state 0. In the preferred form, the class codes are concatenated
9 (e.g., 010110011101) to form an index into the next state table, thus yielding the
10 proper next state. In this manner, the next state table and the corresponding
11 table of pointers, which are addressed by state, is generated for a regular
12 expression. That is, next state table generator 216 works through the state
13 machine and alphabet lookup tables to generate the next state and pointer tables.

14

15 The following is an explanation of the invention from a somewhat different
16 perspective: The purpose of the invention is high speed recognition of patterns in
17 a data stream. The patterns are described by 'regular expressions', which means
18 they may be quite general. For example, the regular expression to detect
19 filenames prefixed by 'binky' or 'winky', containing 'xyz' and having a filename
20 extension '.jpg' are found by the regular expression:

21

```
(binky | winky).*xyz.*\.
```

22 The RDFA (i.e. the present invention) can search for patterns at fixed locations
23 (anchored), as needed for IP packet filtering, but it can also locate unanchored
24 regular expressions anywhere in a packet payload.

25

1 The RDFA has other very important features and advantages over a conventional
2 DFA. It allows parallel processing of bytes. This is important in high speed
3 applications such as OC-192 transport layers, where four bytes arrive from the
4 framer at time. A conventional DFA can not be easily implemented at OC-192
5 rates with todays memory speed cycle time and logic delay time limitations.

6

7 Another advantage is that the RDFA has memory requirements that can be
8 precomputed for a particular set of patterns to be recognized. Finally, the design
9 allows convenient separation of the algorithm between on and off-chip memory
10 when expression complexity becomes large.

11

12 A conventional DFA requires creation of a state machine prior to its use on a data
13 stream and the RDFA has a similar requirement. The user makes a list of regular
14 expressions (rules) and actions to be carried out if a rule is satisfied. A special
15 purpose RDFA compiler converts the rules and actions into a set of tables which
16 are downloaded to the RDFA hardware. The operation of the RDFA is described
17 below, followed by a description of the algorithm implemented in hardware.

18 Finally, the process and algorithms used to created the RDFA tables are
19 described.

20

21 The RDFA operates in a manner similar to a (DFA) but with new features that
22 allow parallelization of the processing, while making enormous reduction in the
23 memory requirements compared with naïve parallelization. The speed of a
24 conventional DFA is often limited by the cycle time of memory used in its
25 implementation. For example, processing the data stream from an OC-192

1 source must handle 10 billion bits/second (10 Gbs). This speed implies a byte
2 must be processed every 0.8 nS, which is beyond the limit of state of the art
3 memory and logic circuits. For comparison, conventional high speed SDRAM
4 chips operate with a 7.5 nS cycle time, which is ten times slower than required. A
5 feature of the RDFA is the ability to process the data stream bytes in parallel. As
6 described below the RDFA processes 4 bytes in parallel, but the algorithm may
7 be applied to arbitrary numbers of bytes, meaning that it is scalable to higher
8 speed. This allows the use of memory that is readily available and far lower in
9 cost than required by a brute-force, single byte at a time, conventional approach.

10
11 In the specific embodiment presented, a separate 256 element lookup table is
12 used for each of the 4 bytes processed in parallel, with table entries having 1 to 4
13 bits width per entry. Fig 3 illustrates the multiple alphabet lookup table concept.
14 In this particular example, the first byte looked up produces a two bit result. The
15 second byte produces a three bit result and so forth. It is very important to note
16 that the set of alphabet lookup tables used are state-dependent. Thus, a different
17 set of four tables is used for each state. Further, the widths of the table entries
18 for each byte position can vary from one state to the next. Thus, the width of
19 each alphabet table lookup is stored as a function of state, so that it can be
20 looked up by state, in order to concatenate the correct number of bits from each
21 lookup. In the Fig 3 example the table widths in bits are denoted (2, 3, 3, 4). The
22 table widths in another state might be (1, 1, 2, 4).

23
24 The lookup results for each byte are concatenated together, which for the Fig 3
25 example produces a 12 bit result. Next, as shown in Fig 4, the current state is

1 used as an index into a table of pointers and a single pointer is selected on the
2 basis of current state. Each pointer, points at a separate linear (1-D) table of next
3 state values and the concatenated result of the parallel alphabet lookup is used
4 as an offset into the selected table. The selected table entry is the value of the
5 next state and also has a code indicating when a terminal state has been
6 reached. The high order bit for each next-state table entry is called the 'special
7 flag' and is set to one to indicate an acceptor state has been reached somewhere
8 in the 4 bytes being processed. In the preferred embodiment, when a lookup in
9 the next-state table results in the special flag being set, the hardware will store
10 two pieces of information into the next entry of a DFA results table. The data
11 saved are:

- 12 1. Pointer to the word (4 bytes) in the packet at which the terminal state
13 occurred.
- 14 2. The table offset (computed from the alphabet table lookups results and
15 index table) into the next-state table.

16

17 Post processing software can use the saved data to determine at which of the 4
18 bytes the regular expression terminated. This is acceptable in situations of
19 interest, only a small number of regular expression matches occur per packet.

20 Alternate embodiments store codes into a secondary terminal state table which
21 will let the hardware directly determine which of the 4 bytes in a word, terminated
22 the matched pattern.

23

24 Regardless of how the special flag is set, the value of the next-state table entry is
25 used to set the next state of the RDFA machine. Thus, when processing the full

1 packet has finished, the entries in the result-table contain information that can be
2 used to determine all regular expression matches in the packet.

3
4 The use of pointers to next state tables, rather than directly using the alphabet
5 table lookup results, allows flexibility in memory management. For example, in
6 embodiments that have on-chip and off-chip memory, pointers can be used so
7 that more frequently used memory is on-chip, to speed up RDFA performance.
8 The expression compiler can determine the amount of memory required. This
9 allows the user to know if a particular set of rules will fit in the on-chip memory.
10 Thus, memory related performance can be accurately known ahead of time.

11
12 The preferred embodiment requires memory lookup operations to process the 4
13 bytes. Specifically, the memory lookups are:

- 14 1. Parallel lookup of each incoming byte from the data stream.
- 15 2. Lookup of the number of bits width for each alphabet lookup result based
16 on state
- 17 3. Lookup of pointer for next-state table based on current state
- 18 4. Lookup of next state.

19 These memory operations may be pipelined to allow effective processing times
20 limited by the longest memory access. Another advantage of the approach is
21 seen when its memory requirements are compared with a simple DFA approach
22 applied to processing 4 bytes in parallel. A simple approach to DFA
23 parallelization, does a lookup on the 4 bytes in parallel. This will match the speed
24 of the RDFA, but requires a table of size 2^{32} entries, which has 4.295 billion
25 entries and a cycle time of 3.2 nS in order to keep up with OC-192 rates (10

1 Gb/sec). Such a system is difficult to implement with current or near-term
2 memory technology, based on the speed and size required. Further, such a large
3 memory is difficult to implement on-chip with the RDFA processing algorithm.

4
5 RDFA Algorithms and Programming (That is, Creation of the RDFA Tables): The
6 front end to the RDFA converts a regular expression to a Nondeterministic Finite
7 Automata (NFA). The NFA is then converted to a Deterministic Finite Automata
8 (DFA). Finally, an optimization is done on the DFA to minimize the number of
9 states. Such processes are well known to compiler writers and are described in
10 the literature. The process described herein converts the state-optimized DFA to
11 an RDFA.

12
13 The RDFA reduces the redundancy in the alphabet representing the input stream
14 by recognizing that at a given state in the DFA the transition decision can be
15 made by grouping letters of the alphabet according to the transition they cause.
16 For example, in the Fig 1 the only symbol which will cause a transition to state 3 is
17 the letter 'n'. From state 2, all other characters transition to the failure state.
18 Thus, for this case, a single bit is sufficient to represent the alphabet because the
19 characters can be mapped to two classes, 'n' and everything else. More
20 complicated transitions can still in general achieve large compression in the
21 number of bits required to represent the alphabet through mapping to character
22 classes. For example, a portion of a regular expression represented as: (a | b | c |
23 g) would be represented in a state transition diagram as shown in Fig. 5. The use
24 of mapping characters to classes changes the 8 bits conventionally used to
25 represent the input character to just a few bits. The basic idea behind mapping

1 characters to classes is recognition that as far as transitions in the state machine
2 are concerned many different characters have identical results. Thus, these
3 many different character codes, which have identical results as far as state
4 transition can be mapped to the same character code, referred to as a 'class'.
5 The alphabet lookup tables perform this mapping function at wire speed. Offline
6 work is required to initially determine them.

7

8 Algorithm for Production of Alphabet Lookup Tables: The list of states that may
9 be reached in a single transition is referred to as '1-closure'. We define the term
10 'n-closure' as the list of states reachable in n transitions from the current state. n-
11 closure is readily calculated recursively as the list of states reachable from n-1
12 closure. In general when M bytes are processed in parallel, a separate alphabet
13 lookup table must be computed for each of the M bytes. Further a different set of
14 tables is computed for each DFA state. Thus if a DFA has L states and M bytes
15 are processed in parallel, a total of $(L \times M)$ alphabet lookup tables must be
16 produced. The algorithm used to produce the M character class tables used to
17 parallel process M bytes, for starting state S, is as follows. The n^{th} lookup table
18 $(1 \leq n \leq M)$ production first computes the n-1 closure. For each symbol in the
19 alphabet, a list of state transitions from the n-1 closure to the next state is made.
20 Now the alphabet mapping is initialized by placing the first symbol in the alphabet
21 into character class 0. The transition list for the next letter in the alphabet is
22 examined and compared with the transitions for the character class 0. If they are
23 identical, then this character is mapped to class 0, otherwise a new class called
24 'class 1' is created and the character is mapped to it. This process proceeds for
25 each character in the character set. If the list of transitions for the character

1 matches an existing class, then it is placed into that class, otherwise it is the first
2 member of a new class. The result of this process is a character class number
3 associated with each symbol in the alphabet. This mapping means that all
4 characters in a given class have an identical list of transitions from the n-1 closure
5 to the next state and thus as far as the state machine behavior is concerned are
6 identical symbols.

7

8 If the total number of classes for a particular lookup table is P . Then, the number
9 of bits, Q , necessary to represent each symbol is $\text{floor}(\log_2 P) + 1$. Q is the width
10 of the table entries in the lookup table. Note that Q is computed for each lookup
11 table separately and varies as a function of both state and byte position. Since
12 the above discussion is abstract, the concept is illustrated below with a simple
13 example for processing 2 bytes in parallel for the portion of a DFA shown in Fig.
14 6A. The illustration will be done only for the lookup tables from the 0 state. The
15 1-closure for this example is, (1, 2, 3, F) where the failure state is denoted by
16 symbol F . The 2-closure is (1, 4, 5, 6, 7, 8, F) in this illustration. Table 1, which is
17 given in an earlier section of this document is a list of non-failure transitions out of
18 state 0 for an alphabet consisting of letters from a to k. Upon inspection, Table 1
19 shows that the alphabet maps to 4 different equivalent characters, meaning that 2
20 bits are sufficient for the table width. Similarly, Table 2 (also given in an earlier
21 section of this document) is list of non-failure state transitions for each symbol for
22 2-closure. In this case the alphabet maps to 8 equivalent characters, so that 3
23 bits are required for the table width. Note that as indicated in the Q value
24 calculation, if the number of equivalent characters had been 5, the table width,
25 would still be 3 bits.

1

2 An important Feature of RDFA: An important property of the RDFA is that the
3 bytes in the data stream are treated as letters in an alphabet and are mapped to
4 character classes. In general, many characters map to a single class, greatly
5 reducing the number of bits necessary to represent an alphabet symbol. As a
6 consequence, when multiple characters are concatenated together and used for a
7 next-state table lookup, the size of the next-state table is greatly reduced, when
8 compared with concatenation of multiple bytes.

9

10 Important Hardware Implementation Feature: The RDFA has many applications,
11 some involving searching full packets for unanchored expressions. The system
12 (i.e. the engine) described above, is well suited to this application. Another
13 application is searching fixed headers for patterns. A special feature incorporated
14 into the RDFA is a programmable datastream bit mask, where each bit
15 corresponds to a sequential word in the input data stream of a packet. For
16 example, an ethernet packet containing 1500 bytes contains 375 words, and a
17 375 bit mask allows complete freedom in selection of words to be processed.
18 When a bit is set on in the data stream mask, the corresponding word is fed to the
19 RDFA. If the bit is turned off then the corresponding word is not seen by the
20 RDFA. This allows a front end filter that operates at line rate which greatly
21 reduces the load on the RDFA when processing fixed position header information.
22 Further, this can lead to reductions in the complexity and memory used by the
23 RDFA. With the above described mask only a small subset of the datastream
24 must be processed and the data that is processed can be handled in a simpler
25 manner, which in turn means larger rule sets can be used for a given amount of

1 memory.

2

3 Reduction of Table Sizes: The RDFA requires a set of alphabet lookup up tables
4 and a next state table for each state. If the number of states can be reduced,
5 then the size of the lookup tables can be reduced. In a classic DFA, when M
6 characters are processed the state machine transitions through M states. For an
7 RDFA it is recognized that processing M bytes in parallel can be treated as a
8 black box, transitioning between two states. For example, as shown in Fig 7A,
9 the character string 'abcdefgh' is intended to be matched. Not counting initial
10 state, a classic DFA has 8 internal states through which it transitions including the
11 acceptor state. However, if 4 bytes are processed in parallel, then only 2 states
12 are needed to represent the transitions as shown in Fig 7B. Note that this is a
13 special case since cyclic graphs, representing wild-cards or arbitrary numbers of
14 character repetitions, may not occur in this type of processing.

15

16 The invention may be embodied in other specific forms without departing from the
17 spirit or central characteristics thereof. While not discussed in detail, incoming
18 data may be evaluated against a plurality of regular expressions simultaneously.
19 In such a case, entering a failure state for one regular expression state machine
20 only terminates processing with respect to that regular expression. The present
21 invention may also be implemented in any of a variety of systems, e.g., to detect
22 a computer virus in e-mail. The present embodiments are therefore to be
23 considered in all respects as illustrative and not restrictive, the scope of the
24 invention being indicated by appending claims rather than by the foregoing
25 description, and all changes that come within the meaning and range of

1 equivalency of the claims are therefore intended to be embraced therein.

2 What is claimed is:

3

1) A system for recognizing a particular pattern formed by a plurality of bytes in a stream of bytes, said system operating in a series of states, and said system including:

- 4 a plurality of look up tables, one for each of said plurality of bytes, each
- 5 look up table having a section associated with each state,
- 6 a first index table which has index data associated with each state, the
- 7 index data associated with each state comprising a plurality of index bits,
- 8 and
- 9 a next state table which has a data word associated with each state, each
- 10 data word having data that specifies the next state and/or a special
- 11 flag,

12 said system performing the following steps in each particular state of a series of

13 states,

- 14 a) interrogating a particular section of each lookup table utilizing the value
- 15 of the associated bytes to determine the value of a series of bits,
- 16 said particular section of each lookup table being determined by the
- 17 particular state,
- 18 b) generating a pointer from a combination of the value of the series of bits
- 19 determined by said lookup tables and the bits in current state
- 20 location in said first index table, said pointer specifying the address
- 21 of a particular data word in said next state table,
- 22 c) reading the particular data word specified by the pointer generated in
- 23 step "b" and based on the contents of said data word proceeding to
- 24 the next state specified by said data word and/or performing special
- 25 operations in response to said flag, and
- 26 d) repeating steps "a" , "b" and "c" until an end of operation is indicated by
- 27 said flag.

2) The system recited in table 1 wherein all of said plurality of bytes are used to interrogate said look up tables in parallel.

1 3) The system recited in claim 1 wherein those bytes that cause the same
2 transition are grouped into classes whereby multiple bytes can be represented by
3 a single class, each class having a single entry in any particular single section of
4 said table.

5
6 4) The system recited in claim 1 wherein at each state, each look up table
7 generates a series of bits, the particular series of bits generated being dependent
8 upon the particular value of the associated of byte.

9
10 5) The system recited in claim 4 wherein said look up tables generate a variable
11 number of bits depending upon the particular state of operation.

12
13 6) The system recited in claim 3 wherein the bits generated by any look up table
14 depends upon the number of character classes of the associated byte.

15
16 7) The system recited in claim 4 wherein the bits generated by said plurality of
17 look up tables are concatenated together and then concatenated with said value
18 from said first index table.

19
20 8) The system recited in claim 4 wherein the bits generated by said plurality of
21 look up tables and said value from said first index table are combined to form the
22 address of a location in said next state table.

23
24 9) The system recited in claim 1 wherein said data stream comprises a stream of
25 bytes that form a data packet, said particular operation is the recognition of a
26 particular pattern of bytes in a data packet and wherein a signal is given when
27 said pattern is recognized.

28
29 10) The system recited in claim 9 wherein when a particular pattern of bytes is
30 recognized, said operation continues to another state to find another occurrence
31 of said pattern of bytes in said packet.

32

33

1 11) The method of recognizing a particular pattern formed by a plurality of bytes
2 in stream of bytes, said method operating in a series of states, and utilizing:

3 a plurality of look up tables, one for each of said plurality of bytes, each

4 look up table having a section associated with each state,

5 a first index table for storing a plurality of index bits associated with each
6 particular state, and

7 a next state table for storing data words, each of which indicates the next
8 state and/or a special operation flag,

9 said method including the following steps in each particular state in a series of
10 states,

11 a) interrogating a particular section of each lookup table utilizing the value
12 of the associated byte to determine the value of a series of bits, said
13 particular section of each look up table being based on the particular
14 state,

15 b) generating a pointer which specifies the address of a data word in said
16 next state table from a combination of the selected values from the
17 lookup tables and a value from the first index table associated with
18 the particular state,

19 c) reading the particular data work specified in step "b" and based on the
20 contents of said data word proceeding to the next state and/or
21 performing operations indicated by a flag in said data work, and

22 d) repeating steps "a" , "b" and "c" until a flag in said next state table
23 indicates an end of the operation.

24
25 12) The method recited in claim 11 wherein all of said plurality of bytes are
26 matched with data in said look up tables in parallel.

27
28 13) The method recited in claim 11 wherein those bytes that cause the same
29 transition are grouped into classes whereby multiple bytes can be represented by
30 a single class, each class having a single entry in a section of said table.

31
32 14) The method recited in claim 11 wherein at each state, each look up table
33 generates a series of bits, the particular series of bits generated being dependent

1 upon the particular value of the associated byte.

2

3 15) The method recited in claim 14 wherein said tables generate a variable
4 number of bits depending upon the particular state of operation.

5

6 16) The method recited in claim 11 wherein the bits generated by any look up
7 table depends upon the class of the associated byte.

8

9 17) The method recited in claim 14 wherein the bits generated by said plurality of
10 look up tables are concatenated together and then concatenated with said value
11 from said first index table.

12

13 18) The method recited in claim 11 wherein said classes are a compressed
14 representation of the alphabet used in said state machine.

15

16 19) The system recited in claim 1 wherein the series of bits generated by each
17 look up table has a number of bits that can vary from state to state.

18

19 20) The method recited in claim 11 wherein the series of bits generated by each
20 look up table has a number of bits that can vary from state to state.

21

22 21) The system recited in claim 1 wherein said system includes four look up
23 tables.

24

25 22) The system recited in claim 1 wherein each of said tables has 256 entries for
26 each state.

27

28 23) The system recited in claim 1 wherein said system includes a number of look
29 up table which is a multiple of four.

30

31 24) The system recited in claim 1 wherein each of said look up tables has a
32 number of entries equal to the values in the alphabet used by said bytes.

33

1 25 An integrated circuit for recognizing a plurality of characters in parallel, said
2 integrated circuit operating in series of states, and said integrated circuit including

3
4 a plurality of look up tables, one for each of said characters, each look up
5 table having a section associated with each state,
6 an index table for storing a plurality of index bits for each state, and
7 a next state table for storing data words, each of which indicates the next
8 state of said state machine,

9 said integrated circuit performing the following steps in a plurality of particular
10 states,

- 11 a) interrogating the section of said look up table associated with said
12 particular state and generating a series of bits based on the value of
13 the associated character,
14 b) generating a pointer which specifies the address of a data word in said
15 next state table from a combination of the selected values from the
16 lookup tables and a value from said first index table,
17 c) reading the particular data word specified in step "b" and proceeding to
18 the state specified by said data word, and
19 d) repeating steps "a" , "b" and "c" until said operation ends.

20
21 26) The integrated circuit recited in claim 25 wherein the number of bits generated
22 by said look up tables and the number of bits in the value from said current state
23 table vary by state, and the total number of bits is the same in all states.

24
25 27) The integrated circuit recited in claim 25 wherein there are four look up tables
26 and four characters are evaluated in parallel.

27
28 28) A method for evaluating in parallel a plurality of characters in a stream of
29 bytes, said method operating in a series of states and utilizing a set of
30 alphabet lookup tables, one lookup table for each of said plurality of
31 characters, each look up table having a section associated with each of
32 said states, and a next state table, said method comprising:

- 1 a) selecting an alphabet lookup table from said set of alphabet tables, as a
2 function of a current state and byte position within said set of bytes,
3 and obtaining a class code from said selected alphabet table
4 corresponding to the character represented by said byte;
- 5 b) selecting a value from said next state table as a function of the current
6 state, and generating an indication of the next state from said class
7 codes of said set of bytes and from said value from said next state
8 table,
- 9 c) repeating steps "a" and "b" until said state machine arrives at
10 acceptance state or until entering a failure state.

- 11
- 12 29) A system for recognizing a particular pattern formed by a number of bytes in
13 a stream of bytes, said number being at least one, said system operating in
14 a series of states, and said system including
15 a look up table for each byte in said number of bytes, each look up table
16 having a section associated with each state,
17 a first index table which has index data associated with each state, the
18 index data associated with each state comprising a plurality of index
19 bits , and
20 a next state table which has a data word associated with each state, each
21 data word having data that specifies the next state and/or a special
22 flag,
23 said system performing the following steps in each particular state of a series of
24 states,
25 a) interrogating a particular section of each lookup table utilizing the value
26 of each of said bytes to determine the value of a series of bits, said
27 particular section of each lookup table being determined by the
28 particular state,

- 1 b) generating a pointer from a combination of the series of bits generated
2 by the lookup tables and bits in the current state location in said first
3 index table, said pointer specifying the address of a particular data
4 word in said next state table,
5 c) reading the particular data word specified by the pointer generated in
6 step "b" and based on the contents of said data word proceeding to
7 the next state specified by said data word and/or performing special
8 operations in response to said flag, and
9 d) repeating steps "a" , "b" and "c" until an end of operation is indicated by
10 said flag.
11
- 12 30) The system recited in claim 1 wherein said system includes a mask with one
13 bit for each of said plurality of bytes in a data packet, and said steps
14 including an initial step of selecting bits for processing utilizing said mask.
15
- 16 31) The method recited in claim 10 wherein said method utilizes a mask with one
17 bit for each of said plurality of bytes in a data packet, and said steps
18 including an initial step of selecting bits for processing utilizing said mask.
19




FIG. 1A (PRIOR ART)




FIG. 1B (PRIOR ART)




FIG. 2A




FIG. 2B




FIG. 2C




FIG. 3




FIG. 4




Figure 4A




FIG. 5

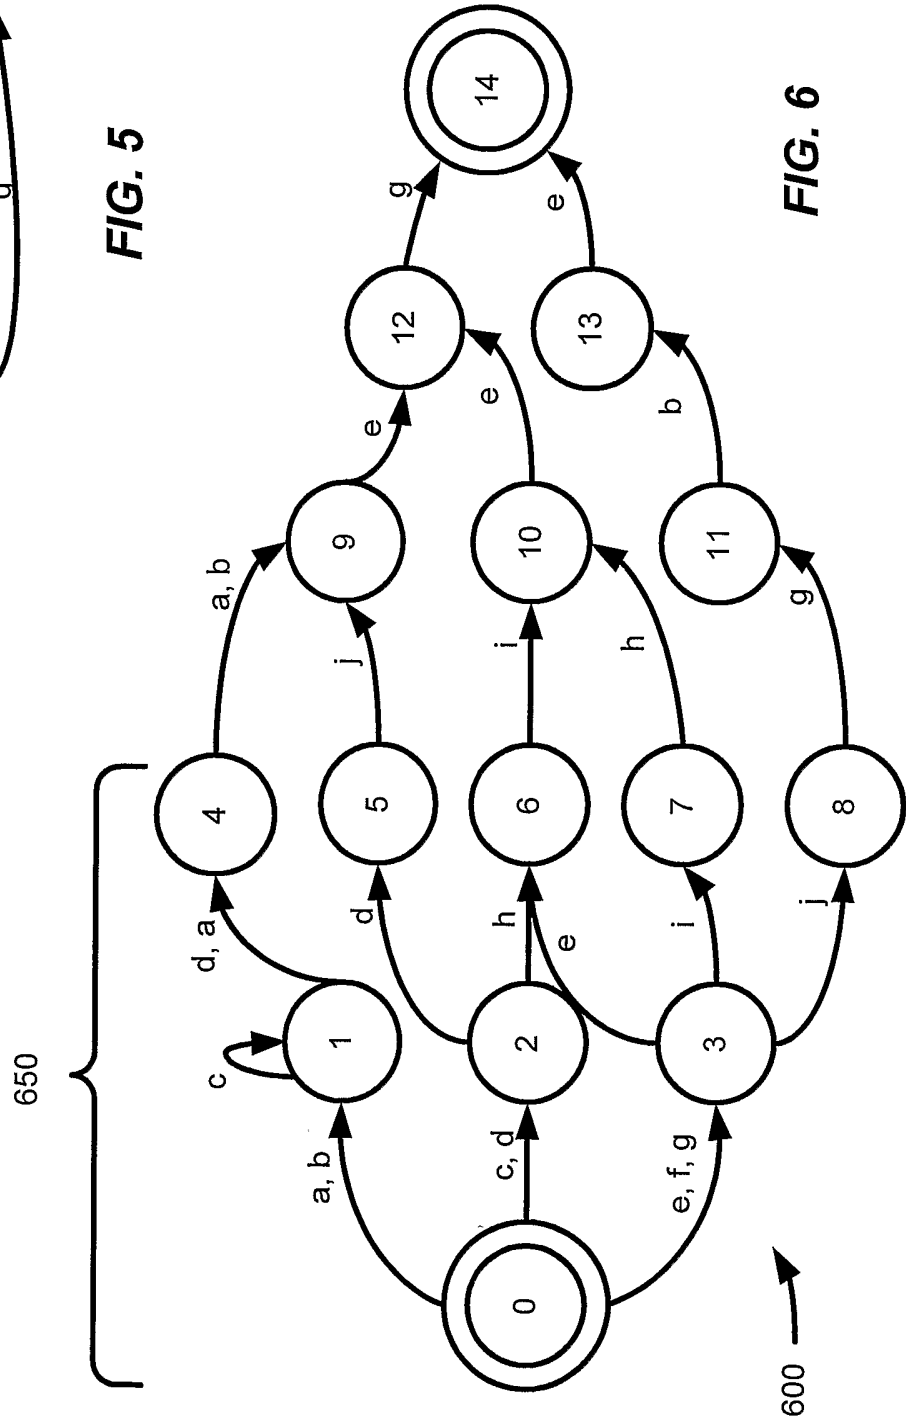


FIG. 6

Fig. 6A Illustration of number of states reachable by 2-closure.




Fig. 7A DFA for processing 8 characters




Fig 7B RDFA for processing 4 bytes in parallel.

