

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
19 February 2004 (19.02.2004)

PCT

(10) International Publication Number
WO 2004/015593 A2

(51) International Patent Classification⁷: G06F 17/30

(21) International Application Number:
PCT/US2003/022574

(22) International Filing Date: 19 July 2003 (19.07.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/215,719 10 August 2002 (10.08.2002) US
10/215,700 10 August 2002 (10.08.2002) US

(71) Applicant: CISCO TECHNOLOGY, INC. [US/US];
Patent Prosecution, c/o Robert Barr, 170 W. Tasman Dr.,
San Jose, CA 95134 (US).

(72) Inventors: EATHERTON, William, N.; 1266 Hickerson
Court, San Jose, CA 95127 (US). LEE, Jaushin; 19869
Lanark Lande, Saratoga, CA 95070 (US). PRIYADAR-
SHAN, Bangalore, L.; 4177 Polonius Circle, Fremont,
CA 94555 (US). WARKHEDE, Priyank, Ramesh;
4355 Renaissance Dr. 201, San Jose, CA 95134 (US).

ERTEMALP, Fusun; 491 Claremont Way, Menlo Park,
CA 94025 (US). HOLBROOK, Hugh, Weber; 85
Buena Vista Terrace, San Francisco, CA 94117 (US).
DEVIREDDY, Dileep, Kumar; 4260 Albany Drive #
I-110, San Jose, CA 95129 (US). KANEKAR, Bhushan,
Mangesh; 4022 Chamberer Drive, San Jose, CA 95135
(US). PULLELA, Venkateshwar, Rao; 3032 Vin Grande
Ct., San Jose, CA 95135 (US).

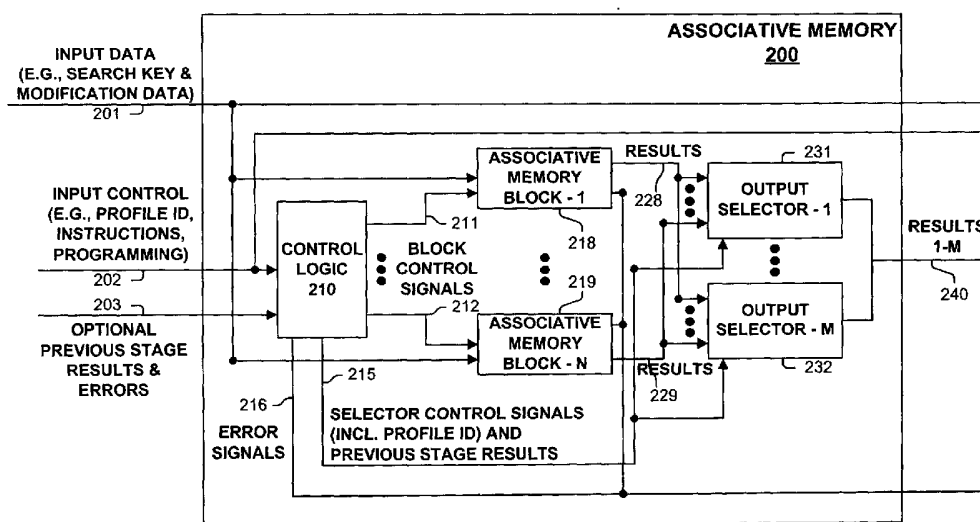
(74) Agent: WILLIAMS, Kirk; The Law Office of Kirk D.
Williams, (USPTO Customer No. 26327), 1234 S. Ogden
St., Denver, CO 80210 (US).

(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, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD,
SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ,
VC, 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),

[Continued on next page]

(54) Title: ASSOCIATIVE MEMORY WITH ENHANCED CAPABILITIES



(57) Abstract: Methods and apparatus are disclosed for performing lookup operations using associative memories, including, but not limited to modifying search keys within an associative memory based on modification mappings, forcing a no hit condition in response to a highest priority matching entry including a force no hit indication, selecting among various sets or banks of associative memory entries in determining a lookup result, and detecting and propagating error conditions. In one implementation, each block retrieves a modification mapping from a local memory and modifies a received search key based on the mapping and received modification data. In one implementation, each of the associative memory entries includes a field for indicating that a successful match on the entry should or should not force a no hit result. In one implementation, an indication of which associative memory blocks or sets of entries to use in a particular lookup operation is retrieved from memory.

WO 2004/015593 A2



European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

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, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, 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, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations 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, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, 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, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA,

- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for all designations
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for all designations

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.

ASSOCIATIVE MEMORY WITH ENHANCED CAPABILITIES

FIELD OF THE INVENTION

This invention especially relates to computer and communications systems, especially network routers and switches; and more particularly, the invention relates to methods and apparatus for performing lookup operations using associative memories, including, but not limited to binary and ternary content-addressable memories, with the lookup operation optionally including modifying a search key in generating a lookup word, forcing no-hit indication in response to matching a particular entry, selectively determining which associative memory blocks to use in identifying a result, and/or propagating error indications.

BACKGROUND OF THE INVENTION

The communications industry is rapidly changing to adjust to emerging technologies and ever increasing customer demand. This customer demand for new applications and increased performance of existing applications is driving communications network and system providers to employ networks and systems having greater speed and capacity (e.g., greater bandwidth). In trying to achieve these goals, a common approach taken by many communications providers is to use packet switching technology. Increasingly, public and private communications networks are being built and expanded using various packet technologies, such as Internet Protocol (IP).

A network device, such as a switch or router, typically receives, processes, and forwards or discards a packet based on one or more criteria, including the type of protocol used by the packet, addresses of the packet (e.g., source, destination, group), and type or quality of service requested. Additionally, one or more security operations are typically performed on each packet. But before these operations can be performed, a packet classification operation must typically be performed on the packet.

Packet classification as required for, *inter alia*, access control lists (ACLs) and forwarding decisions, is a demanding part of switch and router design. The packet classification of a received packet is increasingly becoming more difficult due to ever

increasing packet rates and number of packet classifications. For example, ACLs require matching packets on a subset of fields of the packet flow label, with the semantics of a sequential search through the ACL rules. IP forwarding requires a longest prefix match.

Known approaches of packet classification include using custom
5 application-specific integrated circuits (ASICs), custom circuitry, software or firmware controlled processors, and associative memories, including, but not limited to binary content-addressable memories (binary CAMs) and ternary content-addressable memories (ternary CAMs or TCAMs). Each entry of a binary CAM typically includes a value for matching against, while each TCAM entry typically includes a value and a mask. The
10 associative memory compares a lookup word against all of the entries in parallel, and typically generates an indication of the highest priority entry that matches the lookup word. An entry matches the lookup word in a binary CAM if the lookup word and the entry value are identical, while an entry matches the lookup word in a TCAM if the lookup word and the entry value are identical in the bits that are not indicated by the mask as being irrelevant
15 to the comparison operations.

Associative memories are very useful in performing packet classification operations. In performing a packet classification, it is not uncommon for multiple lookup operations to be performed in parallel or in series using multiple associative memories basically based on a same search key or variant thereof, as one lookup operation might be
20 related to packet forwarding while another related to quality of service determination. Desired are new functionality, features, and mechanisms in associative memories to support packet classification and other applications.

Additionally, as with most any system, errors can occur. For example, array parity errors can occur in certain content-addressable memories as a result of failure-in-time
25 errors which are typical of semiconductor devices. Additionally, communications and other errors can occur. Prior systems are known to detect certain errors and to signal that some error condition has occurred, but are typically lacking in providing enough information to identify and isolate the error. Desired is new functionality for performing error detection and identification.

SUMMARY OF THE INVENTION

Methods and apparatus are disclosed for performing lookup operations using associative memories. In one embodiment, a search key is received and a modification mapping is identified. The search key is modified based on the modification mapping to generate a lookup word. A lookup operation is performed based on the lookup word on multiple associative memory entries. In one embodiment, a profile ID is received and the modification mapping is identified based on the profile ID, such as by retrieving the modification mapping from a memory. In one embodiment, modification data is also received, and the search key is modified by replacing one or more bits or bytes in the search key with data from said modification data based on the modification mapping.

In one embodiment, each of the associative memory entries includes a first field including a value for use in a matching operation on a lookup word and a second field including data for indicating that a successful match on the entry should or should not force a no-hit result. In one embodiment, when a lookup operation is performed and a particular entry including an indication to force a no-hit is identified as a highest-priority matching entry for a set (e.g., one or more entries, an associative memory block, associative memory, etc.) of associative memory entries, the result or intermediate result of the lookup operation for that set is interpreted as a no-hit result (i.e., rather than a successful hit or match on the particular entry) based on the particular entry being identified to a force no-hit condition.

In one embodiment, an indication of a subset of associative memory entries to select among in determining an ultimate lookup result is retrieved from a memory, typically for each output of the associative memory. In one embodiment, indications of identified error conditions are propagated through one or more downstream cascaded associative memories.

BRIEF DESCRIPTION OF THE DRAWINGS

The appended claims set forth the features of the invention with particularity. The invention, together with its advantages, may be best understood from the following detailed description taken in conjunction with the accompanying drawings of which:

5 FIGs. 1A-E are block diagrams of various exemplary systems including one or more embodiments for performing lookup operations using associative memories;

 FIG. 2 is a block diagram of an associative memory including one or more embodiments for performing lookup operations;

10 FIGs. 3A-D illustrate various aspects of a control used in one embodiment for performing lookup operations;

 FIGs. 4A-G illustrate various aspects of an associative memory block used in one embodiment for performing lookup operations; and

 FIGs. 5A-C illustrate various aspects of an output selector used in one embodiment for performing lookup operations.

15

DETAILED DESCRIPTION

Methods and apparatus are disclosed for performing lookup operations using associative memories, including, but not limited to modifying search keys within an associative memory based on modification mappings, forcing a no-hit condition in
5 response to a highest-priority matching entry including a force no-hit indication, selecting among various associative memory blocks or sets or banks of associative memory entries in determining a lookup result, and detecting and propagating error conditions.

Embodiments described herein include various elements and limitations, with no one element or limitation contemplated as being a critical element or limitation. Each of the
10 claims individually recites an aspect of the invention in its entirety. Moreover, some embodiments described may include, but are not limited to, inter alia, systems, networks, integrated circuit chips, embedded processors, ASICs, methods, and computer-readable medium containing instructions. The embodiments described hereinafter embody various aspects and configurations within the scope and spirit of the invention, with the figures
15 illustrating exemplary and non-limiting configurations.

As used herein, the term "packet" refers to packets of all types or any other units of information or data, including, but not limited to, fixed length cells and variable length packets, each of which may or may not be divisible into smaller packets or cells. The term "packet" as used herein also refers to both the packet itself or a packet indication, such as,
20 but not limited to all or part of a packet or packet header, a data structure value, pointer or index, or any other part or identification of a packet. Moreover, these packets may contain one or more types of information, including, but not limited to, voice, data, video, and audio information. The term "item" is used herein to refer to a packet or any other unit or piece of information or data. The phrases "processing a packet" and "packet processing"
25 typically refer to performing some steps or actions based on the packet, and which may or may not include modifying and/or forwarding the packet.

The term "system" is used generically herein to describe any number of components, elements, sub-systems, devices, packet switch elements, packet switches, routers, networks, computer and/or communication devices or mechanisms, or
30 combinations of components thereof. The term "computer" is used generically herein to describe any number of computers, including, but not limited to personal computers,

embedded processing elements and systems, control logic, ASICs, chips, workstations, mainframes, etc. The term "processing element" is used generically herein to describe any type of processing mechanism or device, such as a processor, ASIC, field programmable gate array, computer, etc. The term "device" is used generically herein to describe any type of mechanism, including a computer or system or component thereof. The terms "task" and "process" are used generically herein to describe any type of running program, including, but not limited to a computer process, task, thread, executing application, operating system, user process, device driver, native code, machine or other language, etc., and can be interactive and/or non-interactive, executing locally and/or remotely, executing in foreground and/or background, executing in the user and/or operating system address spaces, a routine of a library and/or standalone application, and is not limited to any particular memory partitioning technique. The steps, connections, and processing of signals and information illustrated in the figures, including, but not limited to any block and flow diagrams and message sequence charts, may be performed in the same or in a different serial or parallel ordering and/or by different components and/or processes, threads, etc., and/or over different connections and be combined with other functions in other embodiments in keeping within the scope and spirit of the invention.

Moreover, the terms "network" and "communications mechanism" are used generically herein to describe one or more networks, communications mediums or communications systems, including, but not limited to the Internet, private or public telephone, cellular, wireless, satellite, cable, local area, metropolitan area and/or wide area networks, a cable, electrical connection, bus, etc., and internal communications mechanisms such as message passing, interprocess communications, shared memory, etc.

The term "storage mechanism" includes any type of memory, storage device or other mechanism for maintaining instructions or data in any format. "Computer-readable medium" is an extensible term including any memory, storage device, storage mechanism, and other storage and signaling mechanisms including interfaces and devices such as network interface cards and buffers therein, as well as any communications devices and signals received and transmitted, and other current and evolving technologies that a computerized system can interpret, receive, and/or transmit. The term "memory" includes any random access memory (RAM), read only memory (ROM), flash memory, integrated

circuits, and/or other memory components or elements. The term "storage device" includes any solid state storage media, disk drives, diskettes, networked services, tape drives, and other storage devices. Memories and storage devices may store computer-executable instructions to be executed by a processing element and/or control logic, and data which is
5 manipulated by a processing element and/or control logic. The term "data structure" is an extensible term referring to any data element, variable, data structure, data base, and/or one or more or an organizational schemes that can be applied to data to facilitate interpreting the data or performing operations on it, such as, but not limited to memory locations or devices, sets, queues, trees, heaps, lists, linked lists, arrays, tables, pointers, etc. A data
10 structure is typically maintained in a storage mechanism. The terms "pointer" and "link" are used generically herein to identify some mechanism for referencing or identifying another element, component, or other entity, and these may include, but are not limited to a reference to a memory or other storage mechanism or location therein, an index in a data structure, a value, etc. The term "associative memory" refers to all types of known or future
15 developed associative memories, including, but not limited to binary and ternary content-addressable memories, hash tables, TRIE and other data structures, etc.

The term "one embodiment" is used herein to reference a particular embodiment, wherein each reference to "one embodiment" may refer to a different embodiment, and the use of the term repeatedly herein in describing associated features, elements and/or
20 limitations does not establish a cumulative set of associated features, elements and/or limitations that each and every embodiment must include, although an embodiment typically may include all these features, elements and/or limitations. In addition, the phrase "means for xxx" typically includes computer-readable medium containing computer-executable instructions for performing xxx.

25 In addition, the terms "first," "second," etc. are typically used herein to denote different units (e.g., a first element, a second element). The use of these terms herein does not necessarily connote an ordering such as one unit or event occurring or coming before the another, but rather provides a mechanism to distinguish between particular units. Additionally, the use of a singular tense of a noun is non-limiting, with its use typically
30 including one or more of the particular item rather than just one (e.g., the use of the word "memory" typically refers to one or more memories without having to specify "memory or

memories," or "one or more memories" or "at least one memory", etc.) Moreover, the phrases "based on x" and "in response to x" are used to indicate a minimum set of items x from which something is derived or caused, wherein "x" is extensible and does not necessarily describe a complete list of items on which the operation is performed, etc.

5 Additionally, the phrase "coupled to" is used to indicate some level of direct or indirect connection between two elements or devices, with the coupling device or devices modify or not modifying the coupled signal or communicated information. The term "subset" is used to indicate a group of all or less than all of the elements of a set. Moreover, the term "or" is used herein to identify an alternative selection of one or more, including all, of the
10 conjunctive items.

Methods and apparatus are disclosed for performing lookup operations using associative memories, including, but not limited to modifying search keys within an associative memory based on modification mappings, forcing a no-hit condition in response to a highest-priority matching entry including a force no-hit indication, selecting
15 among various associative memory blocks or sets or banks of associative memory entries in determining a lookup result, and detecting and propagating error conditions. In one embodiment, each block retrieves a modification mapping from a local memory and modifies a received search key based on the mapping and received modification data. In one embodiment, each of the associative memory entries includes a field for indicating that
20 a successful match on the entry should or should not force a no-hit result. In one embodiment, an indication of which associative memory sets or banks or entries to use in a particular lookup operation is retrieved from a memory.

One embodiment performs error detection and handling by identifying, handling and communication errors, which may include, but is not limited to array parity errors in
25 associative memory entries and communications errors such as protocol errors and interface errors on input ports. Array parity errors can occur as a result of failure-in-time errors which are typical of semiconductor devices. One embodiment includes a mechanism to scan associative memory entries in background, and to identify any detected errors back to a control processor for re-writing or updating the flawed entry. In one embodiment,
30 certain identified errors or received error conditions are of a fatal nature in which no processing should be performed. For example, in one embodiment, a fatal error causes an

abort condition. In response, the device stops an in-progress lookup operation and just forwards error and possibly no-hit signals. Typically, these signals are generated at the time the in-progress lookup operation would have generated its result had it not been aborted so as to maintain timing among devices in a system including the associative memory.

5 In one embodiment, including cascaded or connected associative memory devices, error status messages indicating any error type and its corresponding source are propagated to indicate the error status to the next device and/or a control processor. In addition, the communicated signal may indicate and generate an abort condition in the receiving device. In one embodiment, the receiving device does not perform its next operation or the
10 received instruction, or it may abort its current operation or instruction. Moreover, the receiving device may or may not delay a time amount corresponding to that which its processing would have required in performing or completing the operation or instruction so as to possibly maintain the timing of a transactional sequence of operations.

 FIGs. 1A-E are block diagrams of various exemplary systems and configurations
15 thereof, with these exemplary systems including one or more embodiments for performing lookup operations using associative memories. First, FIG. 1 illustrates one embodiment of a system, which may be part of a router or other communications or computer system, for performing lookup operations to produce results which can be used in the processing of packets. In one embodiment, control logic 110, via signals 111, programs and updates
20 associative memory or memories 115, such as, but not limited to one or more associative memory devices, banks, and/or sets of associative memory entries which may or may not be part of the same associative memory device and/or bank. In one embodiment, control logic 110 also programs memory 120 via signals 123. In one embodiment, control logic 110 includes custom circuitry, such as, but not limited to discrete circuitry, ASICs, memory
25 devices, processors, etc.

 In one embodiment, packets 101 are received by packet processor 105. In addition to other operations (e.g., packet routing, security, etc.), packet processor 105 typically generates one or more items, including, but not limited to one or more packet flow
30 identifiers based on one or more fields of one or more of the received packets 101 and possibly from information stored in data structures or acquired from other sources. Packet processor 105 typically generates a lookup value 103 which is provided to control logic

110 for providing control and data information (e.g., lookup words, modification data, profile IDs, etc.) to associative memory or memories 115, which perform lookup operations and generate one or more results 117. In one embodiment, a result 117 is used is by memory 120 to produce a result 125. Control logic 110 then relays result 107, based on
5 result 117 and/or result 125, to packet processor 105. In response, one or more of the received packets are manipulated and forwarded by packet processor 105 as indicated by packets 109. Note, results 117, 125 and 107 may include indications of error conditions.

FIG. 1B illustrates one embodiment for performing lookup operations using associative memories, including, but not limited to modifying search keys within an
10 associative memory based on modification mappings, forcing a no-hit condition in response to a highest-priority matching entry including a force no-hit indication, selecting among various associative memory blocks or sets or banks of associative memory entries in determining a lookup result, and detecting and propagating error conditions. Control logic 130, via signals 132, programs associative memory or memories 136. In addition,
15 control logic 130 provides control and data information (e.g., lookup words, modification data, profile IDs, etc.) to associative memory or memories 136, which perform lookup operations to generate results and error signals 134, which are received by control logic 130.

FIG. 1C illustrates one embodiment for performing lookup operations using
20 associative memories, including, but not limited to modifying search keys within an associative memory based on modification mappings, forcing a no-hit condition in response to a highest-priority matching entry including a force no-hit indication, selecting among various associative memory blocks or sets or banks of associative memory entries in determining a lookup result, and detecting and propagating error conditions. Control
25 logic 140, via signals 141-143, programs associative memories 146-148. In addition, control logic 140 provides control and data information (e.g., lookup words, modification data, profile IDs, etc.) to associative memories 146-148, which perform lookup operations to generate results and error signals 144-145. As shown each progressive stage forwards error messages to a next associative memory stage or to control logic 140. For example,
30 associative memory 148 relays received error indications via signals 144 via signals 145 to control logic 140.

Moreover, in one embodiment, a synchronization bit field is included in messages 141-145 sent between devices 140 and 146-148, with the value being set or changed at predetermined periodic intervals such that each device 140, 146-148 expects the change. One embodiment uses a single synchronization bit, and if this bit is set in the request or input data 141-145 to a device 146-148, then the device 146-148 will set this bit in the corresponding reply or output data 143-145. For example, in one embodiment, control processor or logic 140 sets the sync bit in its request data 141 periodically, say once in every eight requests. Control processor or logic 140 also monitors the sync bit in the reply data 145. If any kind of error altered the request-reply association (or transaction timing) between the control processor or logic 140 and the associative memories 146-148, then control processor or logic 140 can detect it and recover from that error (by flushing the pipeline, etc.)

In this manner, devices, especially those as part of a transactional sequence, can synchronize themselves with each other. Resynchronization of devices may become important, for example, should an error condition occur, such as an undetected parity error in a communicated instruction signal (e.g., the number of parity errors exceed the error detection mechanism). There is a possibility that a parity error in an instruction goes undetected and that completely changes the transaction timing. Also, there could be other types of "unknown" errors that can put the control processor or logic and the associative memory chain out of synchronization.

FIG. 1D illustrates one embodiment for performing lookup operations using associative memories, including, but not limited to modifying search keys within an associative memory based on modification mappings, forcing a no-hit condition in response to a highest-priority matching entry including a force no-hit indication, selecting among various associative memory blocks or sets or banks of associative memory entries in determining a lookup result, and detecting and propagating error conditions. Control logic 150, via signals 151-153, programs associative memories 156-158. In addition, control logic 150 provides control and data information (e.g., lookup words, modification data, profile IDs, etc.) to associative memories 156-158, which perform lookup operations to generate results and error signals 154-155 which are communicated to control logic 150.

FIG. 1E illustrates a system 180, which may be part of a router or other communications or computer system, used in one embodiment for distributing entries among associative memory units and selectively enabling less than all of the associative memory units when performing a lookup operation. In one embodiment, system 180 includes a processing element 181, memory 182, storage devices 183, one or more associative memories 184, and an interface 185 for connecting to other devices, which are coupled via one or more communications mechanisms 189 (shown as a bus for illustrative purposes).

Various embodiments of system 180 may include more or less elements. The operation of system 180 is typically controlled by processing element 181 using memory 182 and storage devices 183 to perform one or more tasks or processes, such as programming and performing lookup operations using associative memory or memories 184. Memory 182 is one type of computer-readable medium, and typically comprises random access memory (RAM), read only memory (ROM), flash memory, integrated circuits, and/or other memory components. Memory 182 typically stores computer-executable instructions to be executed by processing element 181 and/or data which is manipulated by processing element 181 for implementing functionality in accordance with one embodiment of the invention. Storage devices 183 are another type of computer-readable medium, and typically comprise solid state storage media, disk drives, diskettes, networked services, tape drives, and other storage devices. Storage devices 183 typically store computer-executable instructions to be executed by processing element 181 and/or data which is manipulated by processing element 181 for implementing functionality in accordance with one embodiment of the invention.

In one embodiment, processing element 181 provides control and data information (e.g., lookup words, modification data, profile IDs, etc.) to associative memory or memories 184, which perform lookup operations to generate lookup results and possibly error indications, which are received and used by processing element 181 and/or communicated to other devices via interface 185.

FIG. 2 illustrates an associative memory 200 used in one embodiment for performing lookup operations using associative memories, including, but not limited to modifying search keys within an associative memory based on modification mappings,

forcing a no-hit condition in response to a highest-priority matching entry including a force no-hit indication, selecting among various associative memory blocks or sets or banks of associative memory entries in determining a lookup result, and detecting and propagating error conditions. As shown, control logic 210 receives input control signals 202 which may include programming information. In turn, control logic 210 may update information and data structures within itself, program/update associative memory blocks 218-219, and/or output selectors 231-232. Note, in one embodiment, each of the associative memory blocks 218-219 include one or more associative memory sets or banks of associative memories entries, and logic or circuitry for performing lookup operations.

In one embodiment, input data 201, which may include, but is not limited to search keys and modification data, is received by associative memory 200 and distributed to associative memory blocks 218-219, and possibly forwarded to other downstream associative memories in a cascaded configuration. In addition, input control information 202, which may include, but is not limited to profile IDs (e.g., a value), instructions, programming information, is received by control logic 210, and possibly forwarded to other downstream associative memories in a cascaded configuration. In addition, in one embodiment, previous stage lookup results and/or error indications are received from previous stage associative memories in a cascaded configuration or from other devices by control logic 210. Note, in one embodiment, input data 201, input control 202, previous stage results and errors 203, and/or portions thereof are communicated directly to associative memory blocks 218-219 and/or output selectors 231-232.

Control logic 210 possibly processes and/or forwards the received information via block control signals 211-212 to associative memory blocks 218-219 and via selector control signals and previous stage results 215 (which typically includes the received profile ID) to output selectors 231-232. In addition, control logic 210 may generate error signals 216 based on a detected error in the received information or in response to received error condition indications. Note, in one embodiment, control logic 210 merely splits or regenerates a portion of or the entire received input control 202 and optional previous stage results and errors 203 signals as selector control signals and previous stage results signals 215 and/or error signals 216. In addition, control logic 210 could initiate an abort operation

wherein a lookup operation will not occur because of a detected or received notification of an error condition.

In one embodiment, control logic 210 identifies data representing which associative memory blocks 218-219 to enable, which associative memory blocks 218-219 each output selector 231-232 should consider in determining its lookup result, and/or modification mappings each associative memory block 218-219 should use in modifying an input search key. In one embodiment, this data is retrieved, based on received input control information 202 (e.g., a profile ID or other indication), from one or more memories, data structures, and/or other storage mechanisms. This information is then communicated as appropriate to associative memory blocks 218-219 via block control signals 211-212, and/or output selectors 231-232 via selector control signals and previous stage results signals 215.

In one embodiment, associative memory blocks 218-219 each receive a search key and possibly modification data via signal 201, and possibly control information via block control signals 211-212. Each enabled associative memory block 218-219 then performs a lookup operation based on the received search key, which may include generating a lookup word by modifying certain portions of the search key based on received modification data and/or modification mappings. Each associative memory 218-219 typically generates a result 228-229 which are each communicated to each of the output selectors 231-232. In one embodiment, each associative memory block 218-219 that is not enabled generates a no-hit signal as its corresponding result 228-229. In one embodiment, output selectors 231-232 receive an indication of the associative memory blocks 218-219 that is not enabled.

Output selectors 231 evaluate associative memory results 228-229 to produce results 240. In one embodiment, each output selector has a corresponding identified static or dynamic subset of the associate memory results 228-229 to evaluate in determining results 240. In one embodiment, an identification of this corresponding subset is provided to each output selector 231-232 via selector control signals 215. In one embodiment, each of the output selectors 231-232 receives a profile ID via selector control signals 215 and performs a memory lookup operation based on the received profile ID to retrieve an indication of the particular associate memory results 228-229 to evaluate in determining results 240.

Moreover, in one embodiment, results 240 are exported over one or more output buses 240, each typically connected to a different set of one or more pins of a chip of the associative memory. In one embodiment, the number of output buses used and their connectivity to outputs selectors 231-232 are static, while in one embodiment the number of output buses used and their connectivity to outputs selectors 231-232 are configurable, for example, at initialization or on a per or multiple lookup basis. In one embodiment, an output bus indication is received by an output selector 231-232, which uses the output bus indication to determine which output bus or buses to use. For example, this determination could include, but is not limited to a direct interpretation of the received output bus indication, performing a memory read operation based on the received output bus indication, etc. In one embodiment, an output selector 231-232 performs a memory access operation based on a profile ID to determine which output bus or buses to use for a particular lookup operation. Thus, depending on the configuration, a single or multiple output buses / pins can selectively be used to communicate results 240, with this decision possibly being made based on the tradeoff of receiving multiple results simultaneously versus the number of pins required.

Associative memory 200 provides many powerful capabilities for simultaneously producing one or more results 240. For example, in one embodiment, based on a received profile ID, control logic 210 identifies which of the one or more associative memory blocks 218-219 to enable and then enables them, and provides the profile ID to output selectors 231 for selecting a lookup result among the multiple associative memory blocks 218-219. Each of the associative memory blocks 218-219 may receive/identify a modification mapping based on the profile ID, with this modification mapping possibly being unique to itself. This modification mapping can then be used in connection with received modification data to change a portion of a received search key to produce the actual lookup word to be used in the lookup operation. Also, certain entries may be programmed with force no-hit indications to generate a no-hit result for the corresponding associative memory block 218-219 should a corresponding entry be identified as the highest priority entry matching the lookup word. Each of these enabled associative memories 218-219 typically generate a result (e.g., no-hit, hit with highest priority matching entry or location thereof identified) which is typically communicated to each of the output selectors

231-232. Note, in one embodiment, the results are only communicated to the particular output selectors 231-232 which are to consider the particular result in selecting their respective highest priority result received from associative memory blocks 218-219 and possibly other lookup results from previous stage associative memories. Additionally, in
5 certain configurations, multiple associative memories 200 are cascaded or coupled in other methods so that results from one or more stages may depend on previous stage results, such that a lookup can be programmed to be performed across multiple associative memories 200. These and other constructs provided by associative memory 200 and configurations thereof provide powerful programmable lookup search capabilities and result selection
10 mechanisms using one or more stages of associative memories 200, each including N associative memories blocks 218-219 and M output selectors 231-232. In one embodiment, the actual values of N and M may vary among associative memories 200.

FIG. 3A illustrates a control 300 (which may or may not correspond to control logic 210 of FIG. 2) of an associative memory used in one embodiment. As shown, control 300
15 includes control logic 310 and memory 311. In one embodiment, programming signals 303 are received, and in response, one or more data structures in memory 311 are updated. In addition, control logic generates programming signals 318. In one embodiment, programming 318 is the same as programming signals 303 and thus a physical connection can be used rather than passing through control logic 310. One embodiment of a
20 programming process is illustrated in FIG. 3C, in which processing begins with process block 380. Processing then proceeds to process block 382, wherein programming signals are received. Next, in process block 384, data structures and other elements (e.g., associative memory blocks, output selectors, etc.) are updated. Processing is completed as indicated by process block 386.

25 Returning to FIG. 3A, in performing a lookup operation, input data 301, input control 302, and optionally previous stage results and errors 304 (such as in a cascaded associative memory configuration) are received by control logic 310. In response, one or more data structures in memory 311 are referenced. Control logic 310 generates input data 314, block control signals 315, output selector control signals and (optionally) previous
30 stage results 316, and possibly an error signal 319 indicating a detected error condition or a

received error indicator. In one embodiment, input data 314 is the same as input data 301 and thus a physical connection can be used rather than passing through control logic 310.

FIG. 3B illustrates one set of data structures used in one embodiment. Enable array 320 is programmed with an associative memory block enable indicator 325 for each profile ID 321 to be used. Each associative memory block enable indicator 325 identifies which associative memory blocks are to be enabled for a given lookup operation. Thus, based on a profile ID 321 received via input control 302 (FIG. 3A), enable array 320 can be retrieved from memory 311 (FIG. 3A), which can then be used to generate associative memory block enable signals included in block control signals 315 (FIG. 3A). In one embodiment, associative memory block enable indicator 325 is a bitmap data structure, while in one embodiment, associative memory block enable indicator 325 is a list, set, array, or any other data structure.

Output selector array 330 is programmed with an output selector ID 335 identifying which output selector, such as, but not limited to output selectors 231-232 (FIG. 2) for each tuple (profile ID 331, associative memory block ID 332). Thus, based on a profile ID 331 received over via input control 302 (FIG. 3A), an output selector ID 335 can be identified for each associative memory block ID 332. In one embodiment, output selector ID 335 is a numeric identifier, while in one embodiment, output selector ID 335 is any value or data structure.

Modification mapping array 340 is programmed with a modification mapping 345 for each tuple (profile ID 341, output selector ID 342). Thus, based on a profile ID 341 received over via input control 302 (FIG. 3A), a modification mapping 345 can be identified for each output selector ID 342. In one embodiment, each modification mapping is a data structure identifying how to modify a received search key with received modification data.

FIG. 3D illustrates a process used in one embodiment for initiating a lookup operation. Processing begins with process block 360, and proceeds to process block 362, wherein input data and control signals are received. Next, in process block 364, any previous stage results and error indications are received. As determined in process block 366, if an abort operation should be performed, such as, but not limited to in response to a received fatal error indication or an identified fatal error condition, then processing

proceeds to process block 374 (discussed hereinafter). Otherwise, in process block 368, the enable bitmap, output selector configuration, and modification mappings are received based on the profile ID. Next, in process block 370, data and control signals based on the retrieved and received information are forwarded to the associative memory blocks and
5 output selectors. As determined in process block 372, if an error condition is identified or has been received, then in process block 374, an error indication, typically including an indication of the error type and its source is generated or forwarded. Processing is complete as indicated by process block 376.

FIG. 4A illustrates an associative memory block 400 used in one embodiment.
10 Associative memory block 400 typically includes control logic 410 and associative memory entries, global mask registers, operation logic and priority encoder 412 (e.g., elements for performing the associative memory match operation on a received lookup word). In one embodiment, sets of associative memory entries are grouped into banks of associative memory entries. In one embodiment, programming signals 401 are received,
15 and in response, one or more associative memory entries and/or global mask registers in block 412 are updated. In one embodiment, an associative memory block 400 corresponds to a set or bank of associative memory entries and a mechanism for performing a lookup operation on the set or bank of associative memory entries to produce one or more results. In one embodiment, no mask register is included in associative memory block 400.

20 Moreover, one embodiment of associative memory block 400 includes a memory 413 for storing configuration information, which may allow an associative memory block 400 to retrieve the information from memory 413 rather than receive it from another source. For example, in one embodiment, modification mapping data (e.g., modification mapping 345 of FIG. 3B) or other information is programmed into memory 413. Then,
25 associative memory block 400 retrieves the modification mapping information, such as based on a received profile ID (e.g., rather than receiving the modification mapping signal 404).

Additionally, in one embodiment, a search key 402, modification data 403, modification mapping 404, an enable signal 405, a global mask enable signal 406, and a
30 global mask select signal 407 are received. In response to performing a lookup operation and/or detecting an error condition, such as a parity fault in one of the associative memory

entries, result and error indications 411 are generated. In one embodiment, associative memory entries are checked for parity errors in background. The use of these signals and information in one embodiment are further described in relation to FIGs. 4B-4G.

Turning to FIG. 4B, one embodiment includes multiple global mask registers 415
5 for use in a lookup operation on associative memory entries 416. Global mask enable signal 406 enables the use of a global mask register, while global mask select 407 identifies which of multiple masks to apply to each of the associative memory entries. Lookup word 414 is applied to associative memory entries 416, with possibly using one or more of global masks stored in global mask registers 415, to generate hit/no hit indication 417 and
10 possibly hit location 418 and/or error indication 419, which are incorporated directly or indirectly into result and error indications 411 (FIG. 4A).

FIG. 4C illustrates an error indication 420 used in one embodiment. As shown, error indication 420 includes an error indication 421 for identifying if any or possibly the number of error indications included therein. For any identified error condition or received
15 error indication, an encoded description of each error is included in one or more of the error descriptors 422-423. In one embodiment, a bitmap is used in one or more of error descriptors 422-423, wherein each bit represents a possible error condition, and the value of the bit indicates whether or not a corresponding error has been identified (including received from a prior component or stage.) In one embodiment, each error descriptor
20 422-423 corresponds to a different component, interface, or previous stage. In one embodiment, error indication 420 is used by other components in communicating error conditions or lack thereof.

FIG. 4D illustrates an associative memory entry 430 used in one embodiment. As shown, associative memory entry 430 includes a value 431, an optional mask 432, force no
25 hit indication 433, valid/invalid flag 434, and an error detection value 435. Error detection value 435 may be one or more parity bits, a cyclic redundancy checksum value, or a value corresponding to any other mechanism used for detecting data corruption errors. In one embodiment, value 431 is of a configurable width. In one embodiment, this configurable width includes 80 bits, 160 bits and 320 bits. In one embodiment, such as that of a binary
30 content-addressable memory, no mask field 432 is included. In one embodiment, the width of mask field 432 is variable, and typically, although not required, matches the width of

value field 431. In one embodiment, fields 431-435 are stored in a single physical memory; while in one embodiment, fields 431-435 are stored in multiple physical memories.

FIG. 4E illustrates a mechanism to modify a search key based on modification mapping and modification information used in one embodiment. As shown, a modification mapping bit 443 is used to control selector 440 which selects either search key unit (e.g., one or more bits, bytes, etc.) 441 or modification data unit 442 as the value for lookup unit 445, which is typically a portion of the actual lookup word to be used in matching associative memory entries in a lookup operation.

FIG. 4F illustrates a mechanism to modify a search key 456 based on modification mapping 450 and modification data 454 used in one embodiment. In one embodiment, modification mapping 450 corresponds to a modification mapping 345 (FIG. 3B). As shown in FIG. 4F, modification mapping 450 includes a source portion 451 and a destination portion 452. Referring to the lower portion of FIG. 4F, modification data 454 includes four bytes and search key 456 includes eight bytes. The source portion 451 of modification mapping 450 identifies which bytes of modification data 454 are to be used in generating lookup word 458, and the destination portion 452 of modification mapping 450 identifies where the corresponding bytes to be used of modification data 454 are to be placed in lookup word 458, with the remaining bytes coming from search key 456. In other words, modification mapping 450 and modification data 454 are used to replace certain specified data units in search key 456 in producing the value which will be used in matching the associative memory entries. Of course, various embodiments use different numbers of bits and bytes for modification mapping 450 and modification data 454. In one embodiment, modification mapping 450 includes an indication of the portion of search key 456 to modify (e.g., the value of J in one embodiment, the high-order bytes, the low order bytes, etc.).

FIG. 4G illustrates an associative memory process used in one embodiment in performing a lookup operation. Processing begins with process block 470, and proceeds to process block 472. If the associative memory is not enabled, then processing proceeds to process block 490 wherein a result with a no hit indication is generated, and processing continues to process block 484. Otherwise, in process block 474, the lookup word is determined typically based on the search key, modification mapping, and modification

data. Note, in one embodiment, the search key is used as the lookup word and there is no concept of a modification mapping or modification data. Next, in process block 476, the lookup word is used to match the associative memory entries with consideration of a selected and enabled global mask, if any. Note, in one embodiment, there is no concept of a
5 global mask.

As determined in process block 478, if at least one match has been identified, then processing proceeds to process block 480, otherwise to process block 490, wherein a result with a no hit indication is generated and processing proceeds to process block 484.

Otherwise, as determined in process block 480, if the highest priority matching entry
10 includes a force no hit indication, then processing proceeds to process block 490, wherein a result with a no hit indication is generated and processing proceeds to process block 484. Otherwise, in process block 482, a result indicating a hit (i.e., successful match) with the highest priority matching entry identified is generated.

In process block 484, the result is communicated to at least the identified output
15 selector or selectors. In one embodiment, the output selector to which to communicate the result is identified by output selector ID 335 (FIG. 3B). As determined in process block 486, if an error condition has been identified or received, then in process block 492, a signal is generated indicating the type and location of the error. In one embodiment, error indication 420 (FIG. 4C) is used. Processing is complete as indicated by process block 499.

20 FIG. 5A illustrates of an output selector 500 (which may or may not correspond to an output selector 231-232 of FIG. 2) used in one embodiment. As shown, output selector 500 includes control logic 510 and memory 511. In one embodiment, programming signals 504 are received, and in response, one or more data structures in memory 511 are updated.

FIG. 5B illustrates one data structure used in one embodiment. Available array 520
25 is programmed with an associative memory blocks and optionally previous stage results available for use indicator 525 for each profile ID 521 to be used. Each indicator 525 identifies which, if any, associative memory blocks, sets of entries or associative memory banks are to be considered in determining which matching associative entry to select for the ultimate highest-priority matching associative memory entry. In one embodiment,
30 indicator 525 further identifies which previous stage results to consider. Thus, based on a profile ID 521 received over via selector control signal 501 (FIG. 5A), available array 520

can be retrieved from memory 511 (FIG. 5A). In one embodiment, there is an implied priority ordering of associative memory blocks and any previous stage results, while in one embodiment this priority ordering for determining the ultimate highest-priority matching entry is programmable and/or variable per lookup operation. In one embodiment,
5 associative memory blocks available for use indicator 525 is a bitmap data structure, while in one embodiment, associative memory blocks available for use indicator 525 is a list, set, array, or any other data structure.

Returning to FIG. 5A, in the performance of a lookup operation, output selector 500 receives selector control signal 501, which may include a profile ID. In addition, output
10 selector 500 receives any relevant previous stage results 502 and results 503 from zero or more of the associative memory blocks from which the highest-priority entry will be selected, and which, if any, will be identified in generated result 515.

Moreover, in one embodiment, selector control signal 501 including an enable indication, the enable indication including an enabled or not enabled value, such that in
15 when a not enable value is received, output selector 500 is not enabled and does not select among results from blocks 1-N 503 or optional previous stage results 502. In one embodiment, when not enabled, output selector 500 generates a result signal 515 indicting a no hit, not enabled, or some other predetermined or floating value.

Additionally, in one embodiment, result 515 is communicated over a fixed output
20 bus, which may or may not be multiplexed with other results 515 generated by other output selectors 500. In one embodiment, the associative memory may include one or more output buses, each typically connected to a single pin of a chip of the associative memory, with the selection of a particular output bus possibly being hardwired or configurable, with the configuration possibly being on a per lookup basis, such as that determined from a received
25 value or configuration information retrieved from a memory (e.g., based on the current profile ID.) In such a configuration, control logic 510 (or other mechanism) typically selects which output bus (and the timing of sending result 515) to use for a particular or all results 515.

A process used in one embodiment for receiving and selecting a highest-priority
30 associative memory entry, if any, is illustrated in FIG. 5C. Processing begins with process block 540, and proceeds to process block 542, wherein the results from the associative

memory blocks and the profile ID are received. In process block 544, the set of associative memory blocks to consider in determining the result is retrieved from a data structure/memory based on the profile ID. In process block 546, any relevant previous stage results are received from coupled associative memories. Next, in process block 548, the highest priority match from the available associative memory block and previous stage results is identified, if any. Then, in process block 550, the result is communicated over a fixed or identified output bus/pin or to some other destination, with the result typically including a no hit indication or a hit indication and an identification of the ultimate highest-priority matching associative memory entry. Processing is complete as indicated by process block 552.

In view of the many possible embodiments to which the principles of our invention may be applied, it will be appreciated that the embodiments and aspects thereof described herein with respect to the drawings/figures are only illustrative and should not be taken as limiting the scope of the invention. For example and as would be apparent to one skilled in the art, many of the process block operations can be re-ordered to be performed before, after, or substantially concurrent with other operations. Also, many different forms of data structures could be used in various embodiments. The invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.

20

CLAIMS

What is claimed is:

1. A method for use in performing a lookup operation in an associative memory, the method comprising:

- 5 receiving a search key;
identifying a modification mapping;
modifying the search key based on the modification mapping to generate a lookup word; and
performing a lookup operation on a plurality of associative entries based on the
10 lookup word to generate a result.

2. The method of claim 1, wherein said identifying the modification mapping includes identifying the modification mapping based on a profile ID.

3. The method of claim 2, wherein the modification mapping is retrieved from a memory based on the profile ID.

- 15 4. The method of claim 2, comprising receiving the profile ID.

5. The method of claim 4, comprising receiving modification data; and wherein said modifying the search key includes replacing one or more bits or bytes in the search key with data from said modification data based on the modification mapping.

- 20 6. The method of claim 1, further comprising receiving modification data; and wherein said modifying the search key includes replacing one or more bits or bytes in the search key with data from said modification data based on the modification mapping.

- 25 7. The method of claim 6, wherein said performing a lookup operation includes identifying a particular entry of the plurality of associative entries as a highest priority matching entry; wherein the particular entry includes a force no-hit indication; and wherein the result includes a no-hit value or signal based on the force no-hit indication of the particular entry.

8. The method of claim 1, wherein said performing a lookup operation includes identifying a particular entry of the plurality of associative entries as a highest priority matching entry; wherein the particular entry includes a force no-hit indication; and wherein the result includes a no-hit value or signal based on the force no-hit indication of the particular entry.

9. An associative memory comprising:
a memory for storing a plurality of modification mappings;
a plurality of associative memory entries;
one or more inputs for receiving a search key; and
a control for retrieving a particular modification mapping of said plurality of modification mappings from the memory, for modifying the search key based on the particular modification mapping to generate a lookup word, and for initiating a lookup operation in the plurality of associative memory entries based on the lookup word.

10. The associative memory of claim 9, wherein said one or more inputs is configured for receiving a profile ID; and
wherein the control is configured to retrieve the particular modification mapping from the memory based on the profile ID.

11. The associative memory of claim 10, wherein said one or more inputs is configured for receiving modification data; and
wherein said control is further configured to modify the search key by replacing one or more bits or bytes in the search key with data from said modification data based on the modification mapping.

12. The associative memory of claim 9, wherein said one or more inputs is configured for receiving modification data; and
wherein said control is further configured to modify the search key by replacing one or more bits or bytes in the search key with data from said modification data based on the modification mapping.

13. An associative memory comprising:

means for receiving a search key;

means for identifying a modification mapping;

5 means for modifying the search key based on the modification mapping to generate a lookup word;

means for performing a lookup operation on a plurality of associative entries based on the lookup word to generate a result.

14. The associative memory of claim 13, including means for identifying the modification mapping based on a profile ID.

10 15. The associative memory of claim 14, including means for retrieving the modification mapping from memory based on the profile ID.

16. The associative memory of claim 14, including means for receiving the profile ID.

17. The associative memory of claim 16, comprising:

15 means for receiving modification data; and

means for replacing one or more bits or bytes in the search key with data from said modification data based on the modification mapping.

18. The associative memory of claim 13, comprising:

means for receiving modification data; and

20 means for replacing one or more bits or bytes in the search key with data from said modification data based on the modification mapping.

19. The associative memory of claim 18, wherein said means for performing a lookup operation includes means for identifying a particular entry of the plurality of associative entries as a highest priority matching entry; wherein the particular entry includes a force no-hit indication; and wherein the result includes a no-hit value or signal
25 based on the force no-hit indication of the particular entry.

20. The associative memory of claim 13, wherein said means for performing a lookup operation includes means for identifying a particular entry of the plurality of associative entries as a highest priority matching entry; wherein the particular entry includes a force no-hit indication; and wherein the result includes a no-hit value or signal
5 based on the force no-hit indication of the particular entry.

21. A method for use in performing a lookup operation in an associative memory, the method comprising:

a first associative memory block receiving a profile ID;

10 the first associative memory block retrieving a first modification mapping based on the profile ID;

the first associative memory block receiving a search key and modification data;

the first associative memory block modifying the search key based on the first modification mapping and said modification data to generate a first lookup word;

15 the first associative memory block performing a lookup operation based on the first lookup word to produce a first lookup result;

a second associative memory block receiving the profile ID;

the second associative memory block retrieving a second modification mapping based on the profile ID;

20 the second associative memory block receiving the search key and said modification data;

the second associative memory block modifying the search key based on the second modification mapping and said modification data to generate a second lookup word; and

the second associative memory block performing a lookup operation based on the second lookup word to produce a second lookup result.

25 22. The method of claim 21, wherein the first lookup result has priority over the second lookup result; and the method includes selecting the first lookup result over the second lookup result.

23. The method of claim 21, wherein a particular associative memory entry of the first associative memory block corresponding to the first lookup result includes a force no-hit indication; and wherein the first lookup result includes a no-hit value or signal based on the force no-hit indication of the particular associative memory entry.

5 24. The method of claim 23, wherein the first lookup result has priority over the second lookup result; and the method includes selecting the second lookup result over the first lookup result based on the no-hit value or signal of the first lookup result.

25. An associative memory for producing a first lookup result and a second lookup result, the associative memory comprising:

10 a first associative memory block including:

 a first control,

 a first memory, and

 a first plurality of associative memory entries;

 wherein the first control is configured to receive a profile ID, a search key, and
15 modification data, to retrieve a first modification mapping from the first memory based on the profile ID, to generate a first lookup word based on the search key, said modification data, and the first modification mapping, and to perform a lookup operation on the first plurality of associative memory entries based on the first lookup word to produce the first lookup result; and

20 a second associative memory block including:

 a second control,

 a second memory, and

 a second plurality of associative memory entries;

 wherein the second control is configured to receive the profile ID, the search key,
25 and said modification data, to retrieve a second modification mapping from the second memory based on the profile ID, to generate a second lookup word based on the search key, said modification data, and the second modification mapping, and to perform a lookup operation on the second plurality of associative memory entries based on the second lookup word to produce the second lookup result.

26. The associative memory of claim 25, wherein the first lookup result has priority over the second lookup result; and the associative memory includes means for selecting the first lookup result over the second lookup result.

5 27. The associative memory of claim 25, wherein a particular associative memory entry of the first plurality of associative memory entries corresponding to the first lookup result includes a force no-hit indication; and wherein the first lookup result includes a no-hit value or signal based on the force no-hit indication of the particular associative memory entry.

10 28. The associative memory of claim 27, wherein the first lookup result has priority over the second lookup result; and the associative memory includes means for selecting the second lookup result over the first lookup result based on the no-hit value or signal of the first lookup result.

15 29. A computer-readable medium having stored thereon an associative memory entry, the associative memory entry comprising:
a first field including a value for use in a matching operation on a lookup word; and
a second field including data for indicating that a successful match on the associative memory entry should or should not force a no-hit result.

20 30. The computer-readable medium of claim 29, wherein the associative memory entry comprises a third field including a mask for use in the matching operation on the lookup word.

31. The computer-readable medium of claim 29, wherein the associative memory entry comprises a fourth field indicating whether the entry is valid or not valid.

32. The computer-readable medium of claim 29, wherein the associative memory entry comprises a fifth field indicating an error detection value for the entry.

33. A computer-readable medium having stored thereon an associative memory entry, the associative memory entry comprising:

- a first field including a value for use in a matching operation on a lookup word; and
 - a second field including data to force a no-hit result should the associative memory
- 5 entry be a highest-priority matching entry.

34. A method for generating a lookup result, the method comprising:

- programming a plurality of associative memory entries, at least a particular entry of the plurality of associative memory entries including a force no-hit indication;
 - performing a lookup operation based on a lookup word to identify the particular
- 10 entry as a highest-priority matching entry; and
- generating a no-hit signal in response to the particular entry being identified as the highest-priority matching entry and the particular entry including the force no-hit indication.

35. A method for generating a lookup result, the method comprising:
programming a first plurality of associative memory entries, at least a first particular entry of the first plurality of associative memory entries including a force no-hit indication;

5 programming a second plurality of associative memory entries, the first plurality of associative memory entries having a higher priority than the second plurality of associative memory entries;

performing a lookup operation based on a first lookup word to identify the first particular entry as a first highest-priority matching entry of the first plurality of associative
10 memory entries;

generating a first no-hit signal in response to the first particular entry being identified as the first highest-priority matching entry and the first particular entry including the force no-hit indication;

performing a lookup operation based on a second lookup word to identify the
15 second particular entry as a second highest-priority matching entry of the second plurality of associative memory entries;

generating a second highest-priority signal identifying the second particular entry;
and

identifying the second particular entry as the lookup result in response to the first
20 no-hit signal and the second highest-priority signal.

36. The method of claim 35, wherein a first bank of an associative memory includes the first plurality of associative memory entries, and a second bank of the associative memory includes the second plurality of associative memory entries.

37. An associative memory comprising:
a plurality of associative memory entries, each of the plurality of associative memory entries comprising a value for use in a matching operation on a lookup word and a force or do not force a no-hit indication; and
5 control logic for performing a lookup operation on the plurality of associative memory entries based on a lookup word to identify a particular highest-priority matching entry.

38 The associative memory of claim 37, wherein said control logic signals a no-hit indication when the particular highest-priority matching entry includes a value
10 corresponding to a force a no-hit indication.

39. An apparatus comprising:
means for programming a plurality of entries of an associative memory, at least a particular entry of the plurality of entries including a force no-hit indication;
means for performing a lookup operation based on a lookup word to identify the
15 particular entry as a highest-priority matching entry; and
means for generating a no-hit signal in response to the particular entry being identified as the highest-priority matching entry and the particular entry including the force no-hit indication.

40. The apparatus of claim 39, wherein the associative memory is a ternary or
20 binary content-addressable memory.

41. A method for generating a lookup result, the method comprising:
retrieving an indication of a subset of a plurality of associative memory blocks to use in determining the lookup result, each of the subset of the plurality of associative memory blocks indicating an associative memory block lookup result;
25 selecting a highest priority hit result from said associative memory block lookup results from the subset of the plurality of associative memory blocks; and
indicating the highest priority hit result as the lookup result.

42. The method of claim 41, comprising receiving a profile ID; wherein the indication of the subset of the plurality of associative memory blocks to use is retrieved from a memory based on the profile ID.

43. The method of claim 42, further comprising identifying, based on the profile ID,
5 which of the plurality of associative memory blocks to enable in generating the lookup result.

44. The method of claim 41, wherein said selecting the highest priority hit result further includes selecting from one or more previous stage results.

45. The method of claim 41, wherein each of the plurality of associative memory
10 blocks includes one or more associative memory banks, each of said one or more associative memory banks including a plurality of associative memory entries.

46. The method of claim 41, wherein each of the plurality of associative memory blocks includes a plurality of associative memory entries.

47. The method of claim 41, comprising selecting a particular one of a plurality of
15 output buses, wherein said indicating includes indicating the lookup result on the particular one of the plurality of output buses.

48. The method of claim 47, comprising receiving an indication of the particular one of the output buses; wherein said selecting is performed based on the indication.

49. The method of claim 47, comprising retrieving from a memory an indication of
20 the particular one of the output buses; wherein said selecting is performed based on the indication.

50. The method of claim 49, comprising receiving a profile ID; wherein said retrieving is performed based on the profile ID.

51. An associative memory device comprising:
a plurality of associative memory blocks, each of the associative memory blocks
indicating an associative memory block lookup result; and
an output selector configured to retrieve from a memory an indication of the
5 plurality of associative memory blocks to consider in producing an overall lookup result,
and to select among said plurality of associative memory block lookup results based on the
indication of the plurality of associative memory blocks to consider in producing the
overall lookup result.

52. The associative memory device of claim 51, wherein the associative memory
10 device is configured to receive one or more previous stage results; and wherein the output
selector is configured to further select among said one or more previous stage results in
addition to the plurality of associative memory blocks in producing the overall lookup
result.

53. The associative memory device of claim 52, wherein the indication further
15 indicates which of said one or more previous stage results to select among in producing the
overall lookup result.

54. The associative memory device of claim 53, wherein the associative memory
device is configured to receive a profile ID; and wherein the output selector is configured
to retrieve the indication of the plurality of associative memory blocks to consider from the
20 memory based on the profile ID.

55. The associative memory device of claim 54, wherein the associative memory
device is configured to receive one or more previous stage error signals indicating one or
more previous stage error conditions, and to generate an error signal including said one or
more previous stage error conditions.

25 56. The associative memory device of claim 55, further configured to identify one
or more internal or communications error conditions, and wherein said error signal further
identifies said one or more internal or communications error conditions.

57. The associative memory device of claim 51, wherein the associative memory device is configured to receive a profile ID; and wherein the output selector is configured to retrieve the indication of the plurality of associative memory blocks to consider from the memory based on the profile ID.

5 58. The associative memory device of claim 51, comprising means to identify an enabled subset of the plurality of associative memory blocks to enable for performing a particular lookup operation; and means to communicate an enable indication to the enabled subset of the plurality of associative memory blocks.

10 59. The associative memory device of claim 51, wherein each of the plurality of associative memory blocks includes one or more associative memory banks, each of said one or more associative memory banks including a plurality of associative memory entries.

60. The associative memory device of claim 51, wherein each of the plurality of associative memory blocks includes a plurality of associative memory entries.

15 61. An apparatus for performing lookup operations, the apparatus comprising:
a first associative memory configured to receive a lookup request and to perform a lookup operation to generate a first set of results, the first associative memory including means for identifying a first set of error conditions and means for relaying indications of the first set of results and one or more indications of the first set of error conditions to a second associative memory; and
20 the second associative memory, coupled to the first associative memory, configured to receive said indications of the first set of results and said indications of the first set of error conditions, the second associative memory including means for forwarding said one or more indications of the first set of error conditions to a next device.

25 62. The apparatus of claim 61, wherein the next device includes a third associative memory.

63. The apparatus of claim 61, wherein the first associative memory includes a plurality of associative memory entries, and the first set of error conditions includes a parity error in at least one of the plurality of associative memory entries.

64. The apparatus of claim 61, wherein the first set of error conditions includes a communication error.

65. The apparatus of claim 61, wherein the first set of error conditions includes detection of invalid control information.

5 66. An apparatus for generating a lookup result, the apparatus comprising:
 means for retrieving an indication of a subset of a plurality of associative memory
 blocks to use in determining the lookup result, each of the subset of the plurality of
 associative memory blocks indicating an associative memory block lookup result;
 means for selecting a highest priority hit result from said associative memory block
10 lookup results from the subset of the plurality of associative memory blocks; and
 means for indicating the highest priority hit result as the lookup result.

67. The apparatus of claim 66, comprising means for receiving a profile ID; and
means for identifying, based on the profile ID, the indication of the subset of the plurality
of associative memory blocks to use.

15 68. The apparatus of claim 67, further comprising means for identifying, based on
the profile ID, which of the plurality of associative memory blocks to enable in generating
the lookup result.

69. The apparatus of claim 66, wherein said means for selecting the highest priority
hit result includes means for selecting from one or more previous stage results.

20 70. The apparatus of claim 66, wherein each of the plurality of associative memory
blocks includes one or more associative memory banks, each of said one or more
associative memory banks including a plurality of associative memory entries.

71. The apparatus of claim 66, wherein each of the plurality of associative memory
blocks includes a plurality of associative memory entries.

25 72. The apparatus of claim 66, comprising means for identifying and reacting to an
abort condition.

73. The apparatus of claim 66, comprising means for receiving one or more error signals from one or more previous stage results.

74. The apparatus of claim 73, comprising means for identifying and reacting to an abort condition.

5 75 The apparatus of claim 66, comprising means for communicating a means for indicating one or more error conditions to another device or component.

76. An apparatus for performing lookup operations and to simultaneously produce one or more lookup results, the apparatus comprising:

10 a plurality of associative memory blocks, each of the associative memory blocks indicating an associative memory block lookup result; and
 a plurality of output selectors, each of the plurality of output selectors configured to retrieve from a memory an indication of the plurality of associative memory blocks to select among in producing a selected lookup result, and configured to select among said plurality of associative memory block lookup results based on the indication of the
15 plurality of associative memory blocks to select among in producing the selected lookup result.

77. The apparatus of claim 76, wherein each of the plurality of output selectors is configured to receive one or more previous stage results and to further select among said one or more previous stage results in addition to the plurality of associative memory blocks
20 in producing the selected lookup result.

78. The apparatus of claim 76, comprising means for receiving a profile ID and means for communicating the profile ID to each of the plurality of output selectors; and wherein each of the plurality of output selectors is configured to retrieve the indication of the plurality of associative memory blocks to select among from the memory based on the
25 profile ID.

79. The apparatus of claim 76, wherein the indication corresponding to a particular one of the plurality output selectors indicates to select none of the plurality of associative memory blocks.

80. The apparatus of claim 76, wherein each of the plurality of output selectors is configured to receive a corresponding enable indication, the enable indication including an enabled or not enabled value, wherein each of the plurality of output selectors is configured not to select among the plurality of associative memory blocks when its said corresponding
5 enable indication includes said not enabled value.

81. The apparatus of claim 76, comprising a plurality of output buses, wherein each of the plurality of output selectors is configured to select a particular output bus of the plurality of output buses, and to signal the selected lookup result over the particular output bus.

10 82. The apparatus of claim 81, comprising means for receiving a profile ID and means for communicating the profile ID to each of the plurality of output selectors; wherein each of the plurality of output selectors is configured to retrieve from the memory an indication of the particular output bus.

15 83. The apparatus of claim 82, wherein the particular output buses selected by at least two of the plurality of output selectors are different output buses of the plurality of output buses.

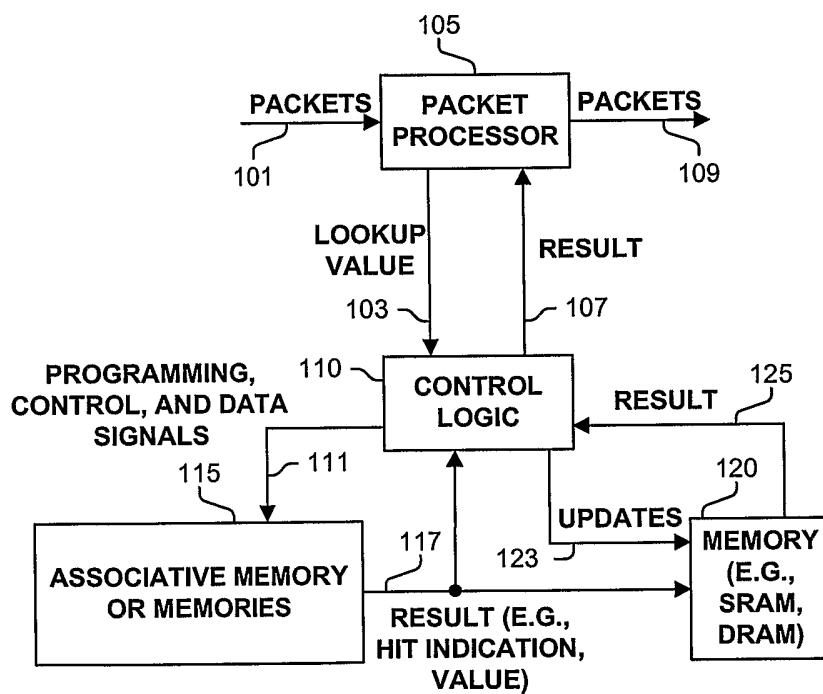


FIGURE 1A

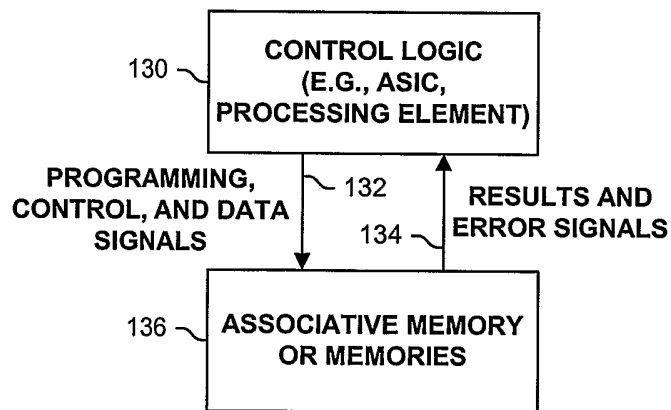


FIGURE 1B

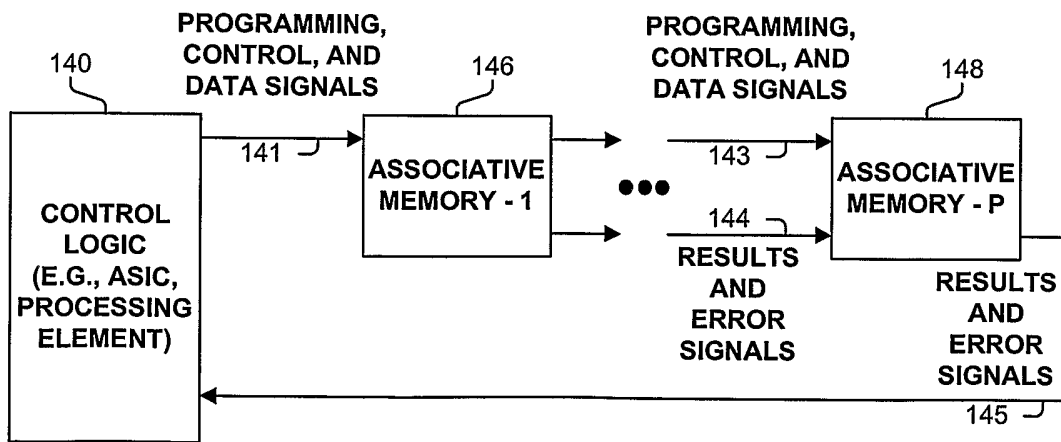


FIGURE 1C

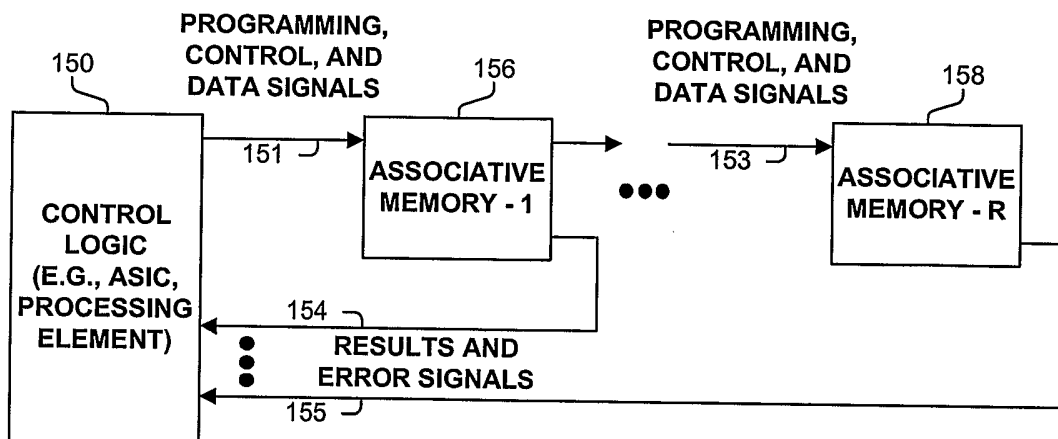


FIGURE 1D

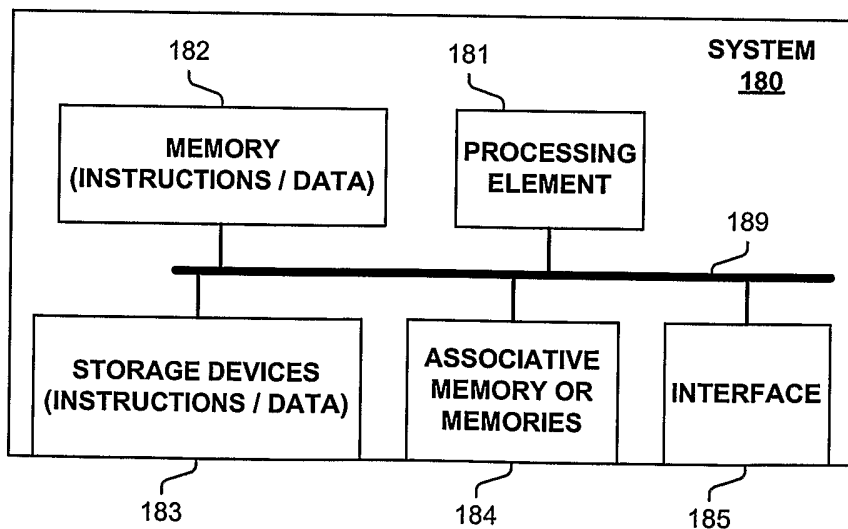


FIGURE 1E




FIGURE 2

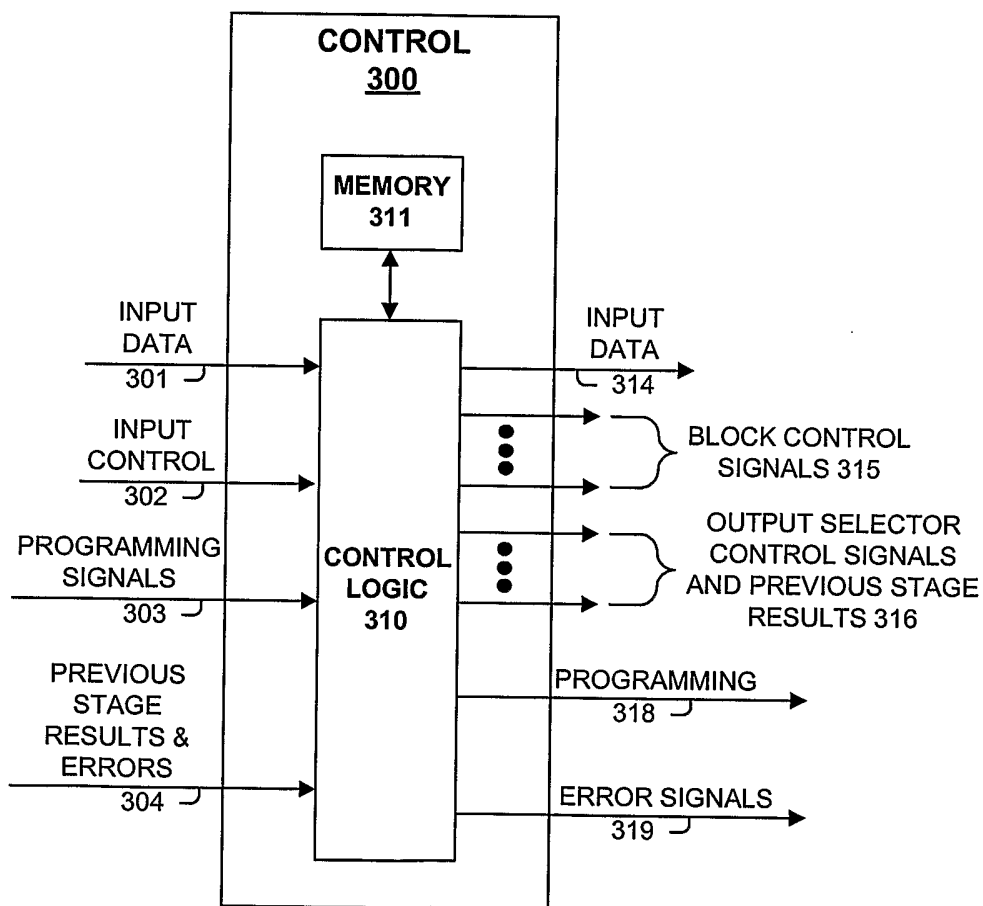


FIGURE 3A

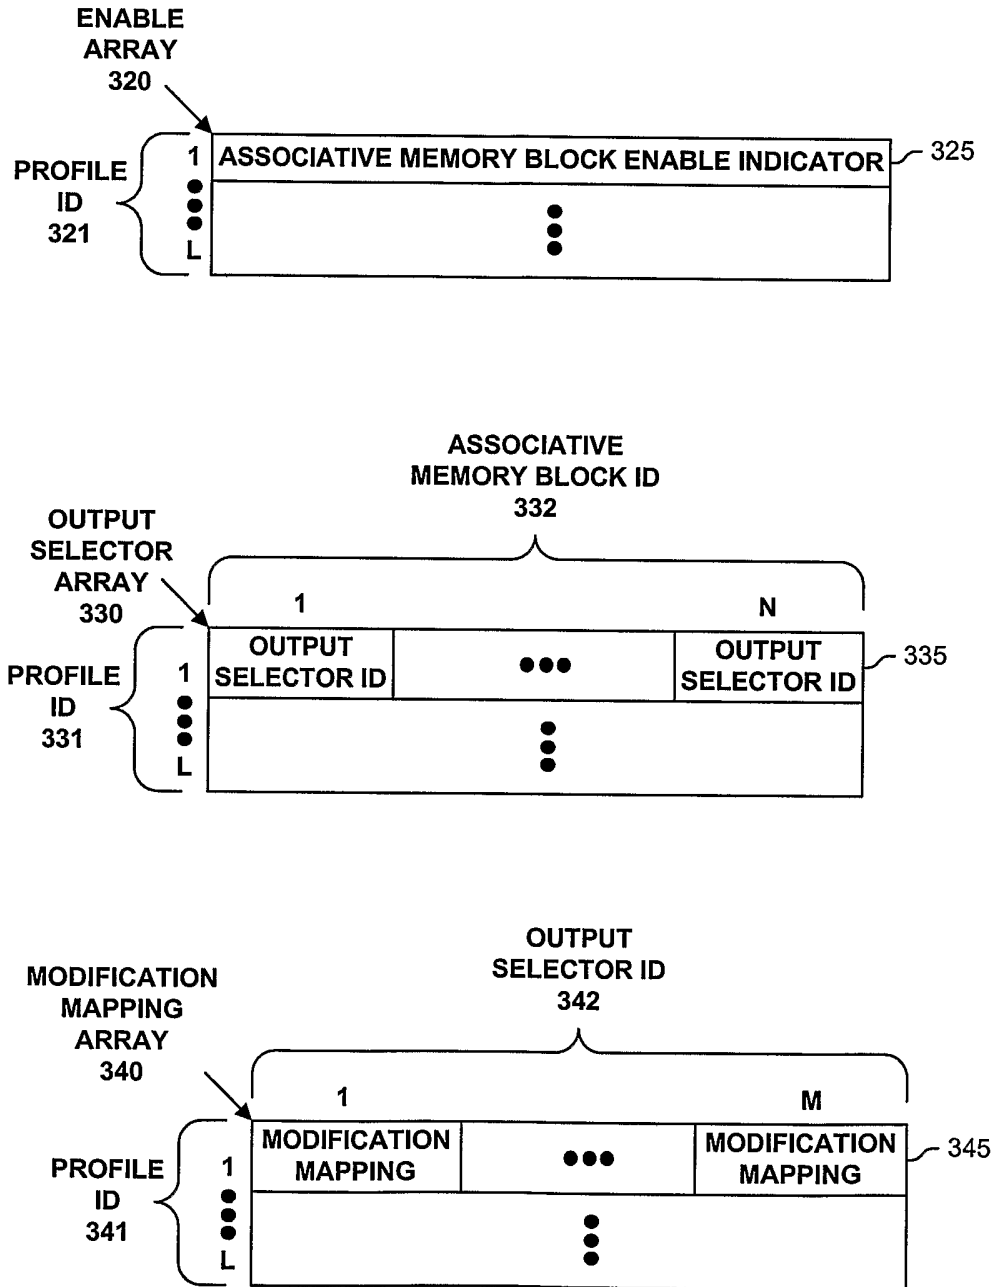


FIGURE 3B




FIGURE 3C

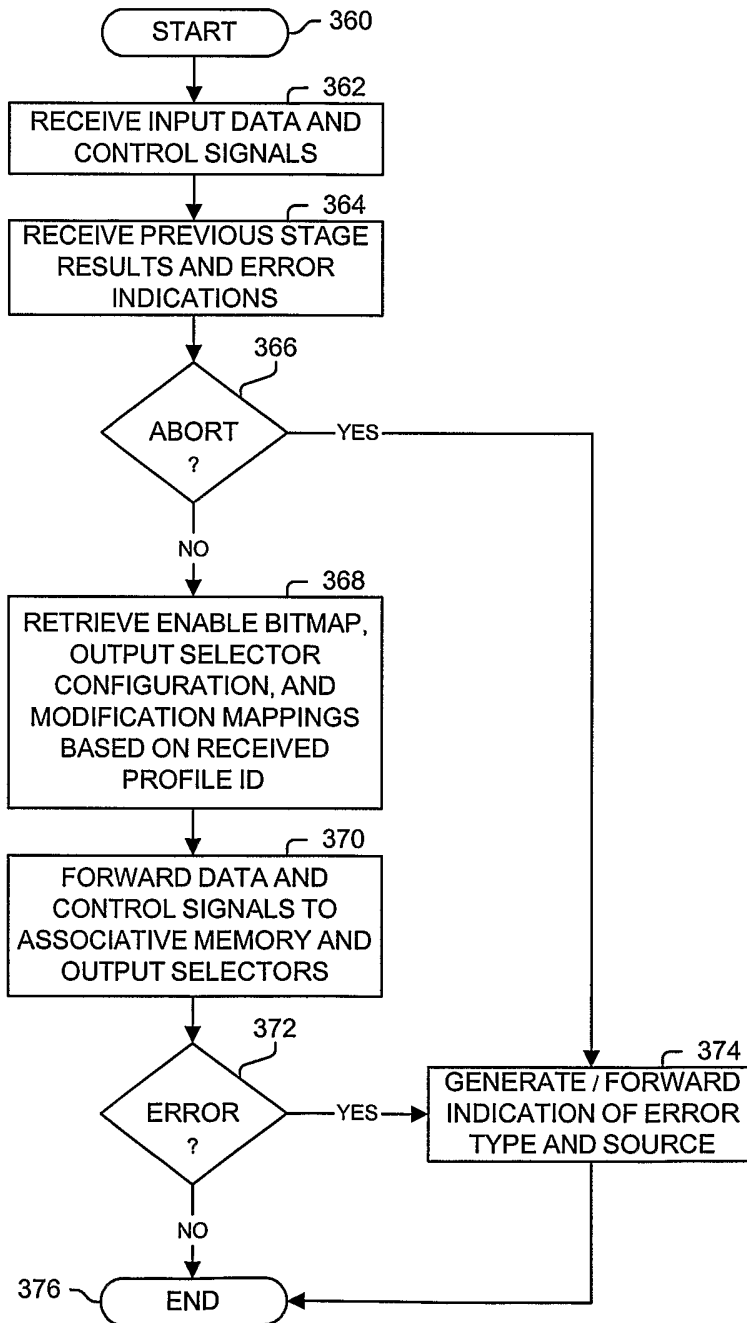


FIGURE 3D

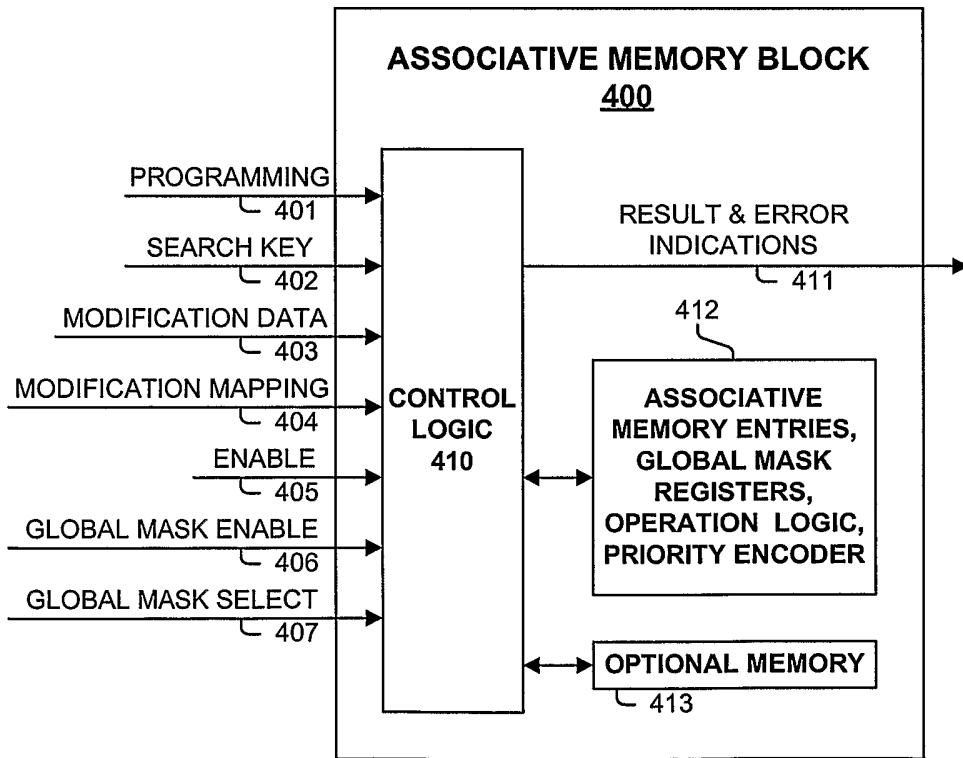


FIGURE 4A

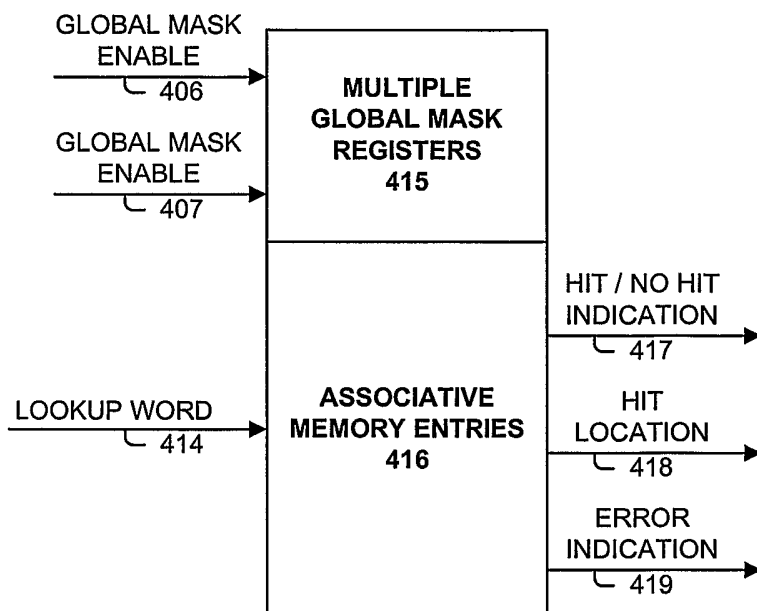


FIGURE 4B

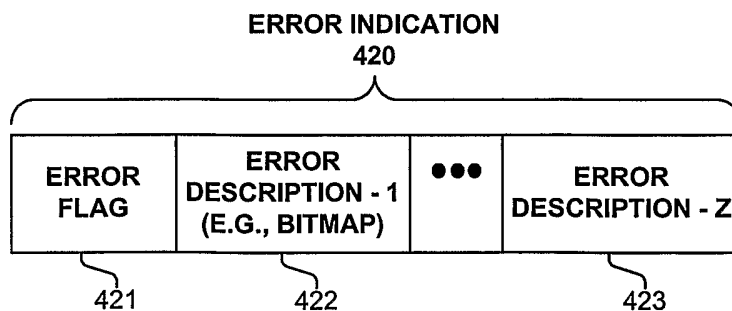


FIGURE 4C




FIGURE 4D




FIGURE 4E

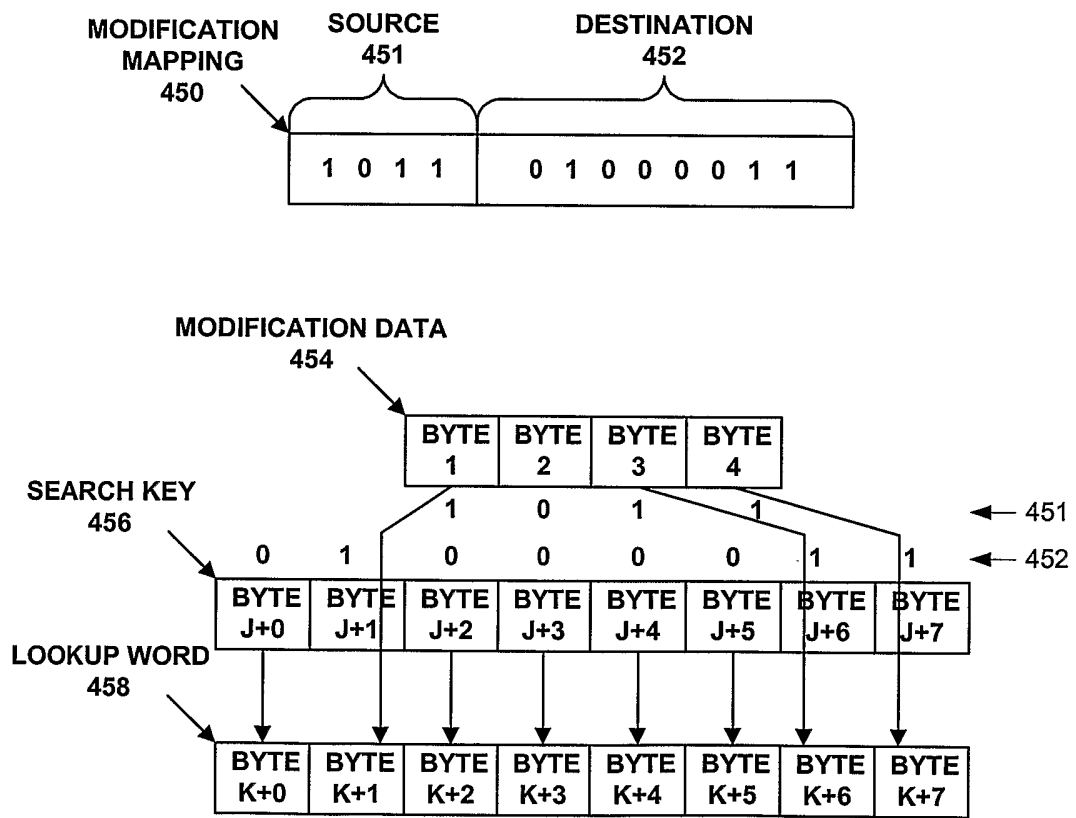


FIGURE 4F

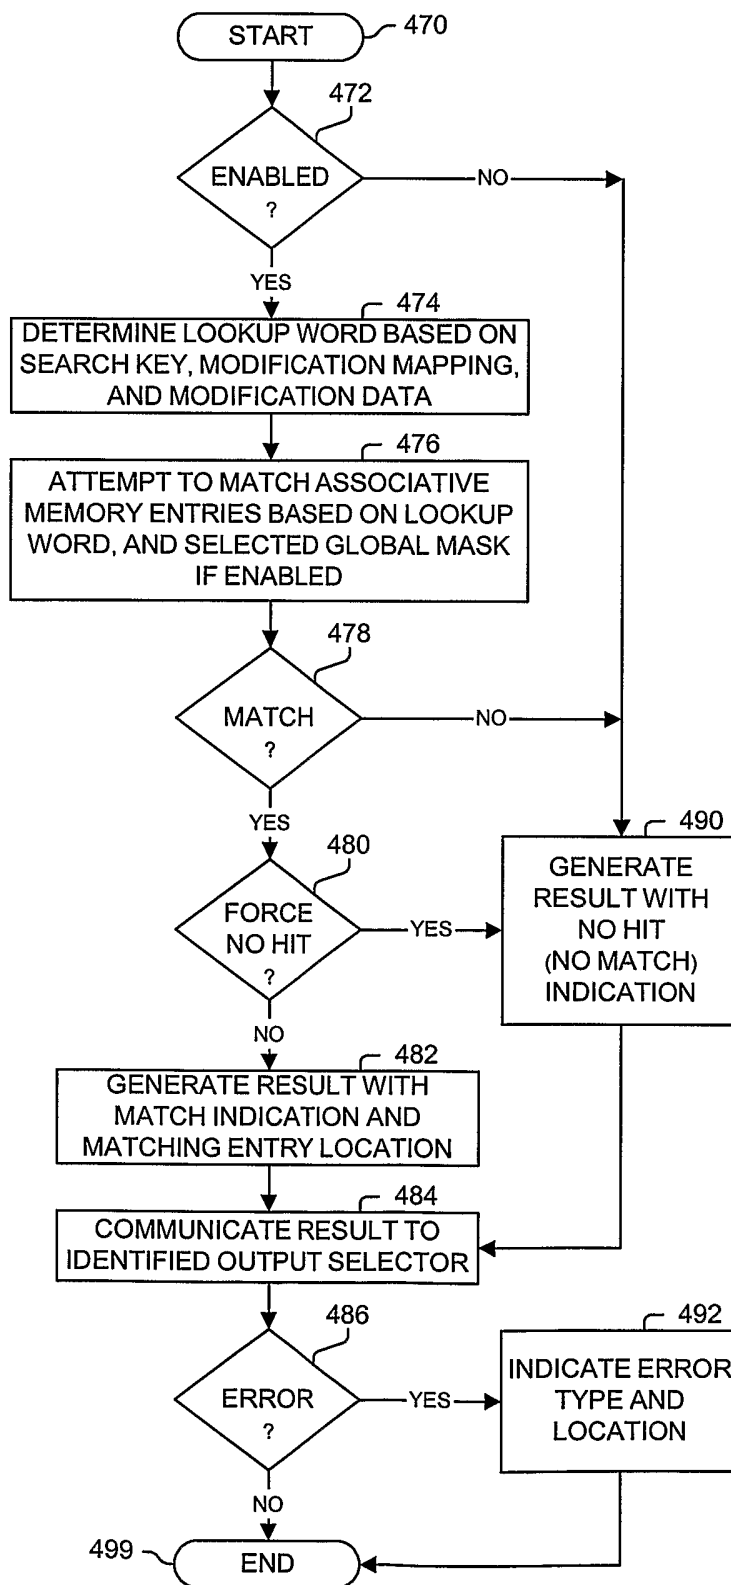


FIGURE 4G

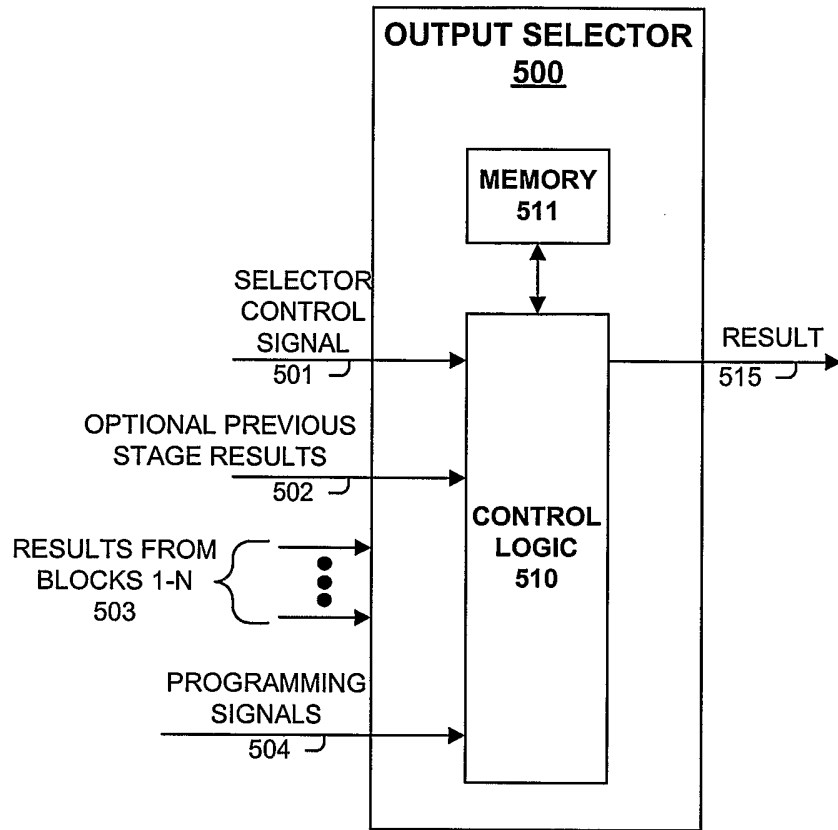


FIGURE 5A

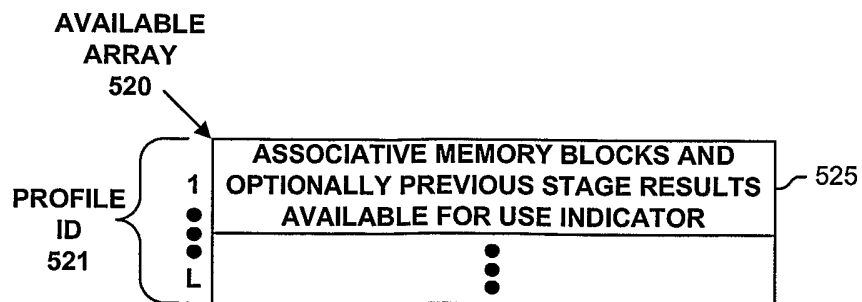


FIGURE 5B

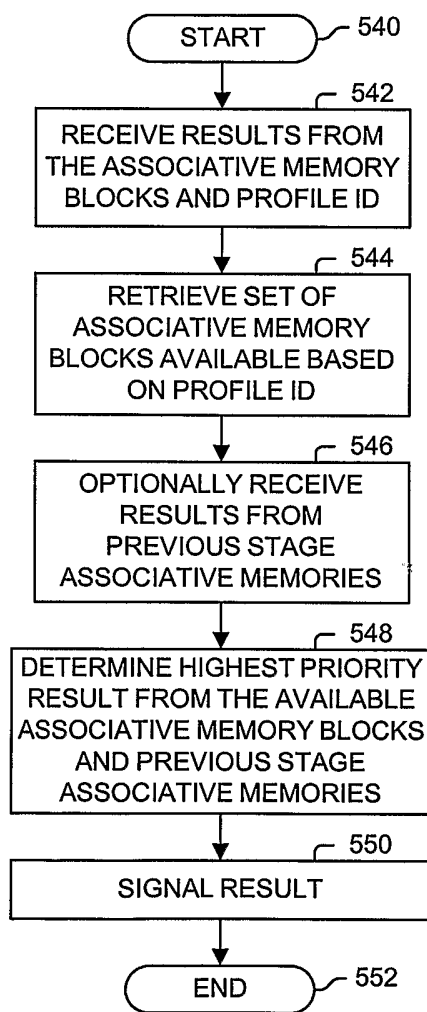


FIGURE 5C