

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
17 September 2009 (17.09.2009)

PCT

(10) International Publication Number
WO 2009/114691 A2

(51) International Patent Classification:
G06F 12/08 (2006.01) *G06F 12/00* (2006.01)

(21) International Application Number:
PCT/US2009/036958

(22) International Filing Date:
12 March 2009 (12.03.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/036,181 13 March 2008 (13.03.2008) US

(71) Applicant (for all designated States except US): **SAND-BRIDGE TECHNOLOGIES, INC.** [US/US]; 120 White Plains Road, 4th Floor, Tarrytown, NY 10591 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **HOANE, Arthur, Joseph**; 10 Oakland Avenue, Yonkers, NY 10710 (US).

(74) Agent: **KARCESKI, Jeffrey, D.**; BARNES & THORNBURG LLP, 750 17th Street, N.W., Suite 900, Washington, DC 20006 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (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 (Rule 48.2(g))



WO 2009/114691 A2

(54) Title: METHOD FOR ACHIEVING POWER SAVINGS BY DISABLING A VALID ARRAY

(57) Abstract: A method is described for use when a cache is accessed. Before all valid array entries are validated, a valid array entry is read when a data array entry is accessed. If the valid array entry is a first array value, access to the cache is treated as being invalid and the data array entry is reloaded. If the valid array entry is a second array value, a tag array entry is compared with an address to determine if the data array entry is valid or invalid. A valid control register contains a first control value before all valid array entries are validated and a second control value after all valid array entries are validated. After the second control value is established, reads of the valid array are disabled and the tag array entry is compared with the address to determine if a data array entry is valid or invalid.

**Method for Achieving Power
Savings by Disabling A Valid Array**

Cross-Reference to Related Application(s)

[0001] This is a PCT Patent Application that relies for priority on U.S. Provisional Patent Application Serial No. 61/036,181, filed on March 13, 2008, the contents of which are incorporated herein by reference.

Field of the Invention

[0002] The invention relates to a method for identifying if selected entries in a cache are valid. More specifically, the invention relates to method, using a valid array, to determine if entries in a cache are valid. After all of the entries in the cache are found to be valid, the valid array is disabled.

Background of the Invention

[0003] One way of implementing a cache relies upon the use of a separate, valid array. A valid array indicates if specific entries within a cache that are valid. Typically, as is known in the prior art, for each access made to a cache, a read is made of the valid array to determine if the cache entry is valid before accepting the entry in the cache. This means that each time the cache is accessed, a separate read of the valid array is required.

[0004] As should be apparent, however, once all of the entries in a cache have been validated (i.e., all of the entries in the cache have been determined to be valid), the valid array no longer serves a purpose. Once each of the entries in the cache is validated, further reads of the valid array do not produce any result other than the already-determined fact that all of the entries are valid. As a result, further reads of the valid array consume power unnecessarily. Any further reads of the valid array are merely duplicative of prior processing steps.

[0005] As may be appreciated by those skilled in the art, reading of the valid array after all of the entries in the cache have been validated adds unnecessarily to the power consumed by the processor. This is considered to be wasteful.

[0006] As should be appreciated by those skilled in the art, the ability to reduce processor time and power consumption is desirable. Less processor time and lower power consumption contribute to a more efficient processor.

Summary of the Invention

[0007] It is, therefore, one aspect of the invention to improve the efficiency of a processor.

[0008] It is a further aspect of the invention to reduce the power consumed by a processor by reducing the number of times that a valid array is accessed after the entries in a cache have been validated.

[0009] One specific aspect of the invention, therefore, is implementation of an instruction to disable reads to a valid array after all of the entries in a cache have been validated.

[0010] By disabling reads of a valid array after entries in a cache have been validated, the invention reduces the power required by a processor that accesses the cache.

[0011] It is an aspect of the invention, therefore, to provide a method and architecture for disabling reads to the valid array under program control.

[0012] A further aspect of the invention is to provide a method and a mechanism by which a program first may ensure that all entries are valid, and then disable the valid array.

[0013] Other aspects of the invention will become apparent from the discussion provided herein.

Brief Description of the Drawings

[0014] Various aspects of the invention will be described in connection with the drawings incorporated herein, in which:

[0015] Fig. 1 is a flow diagram illustrating a first part of one embodiment of the method of the invention;

[0016] Fig. 2 is a continuation of the flow diagram illustrated in Fig. 1, providing a second part of one embodiment of the method of the invention; and

[0017] Fig. 3 is a continuation of the flow diagram illustrated in Fig. 2, providing a third part of one embodiment of the method of the invention.

Description of Embodiment(s) of the Invention

[0018] The invention will now be described in connection with one or more embodiments. The embodiment(s) described are not intended to be limiting of the invention. To the contrary, the embodiment(s) are intended to illustrate the broad scope of the invention. As should be appreciated by those skilled in the art, there are numerous variations and equivalents of the embodiments that may be employed. Those variations and equivalents are intended to fall within the scope of the invention.

[0019] A cache structure in a processor typically has three (3) components: (1) a tag array, (2) a data array, and (3) a valid array. For purposes of this discussion, the term "processor" is intended to refer broadly to any device capable of executing instructions, whether via software or via hardware. Processors, therefore, encompass computers as a common example. As should be appreciated by those skilled in the art, a processor also may be a handheld device such as a persona data assistant, cellular phone, or the like.

[0020] A typical cache may have 2^m entries. Each entry may have a size of 2^b bytes. While these parameters refer to a typical cache, caches with a greater or a fewer number of entries may be employed without departing from the scope of the invention. Moreover, the invention is not intended to be limited solely to caches where each entry has a size of 2^b bytes. The entries may have a larger or a smaller byte sizes without departing from the scope of the invention.

[0021] During a cache access, the tag array and the data array are accessed based on bit $m+b-1$ to bit b of the address. The contents read from the tag-array are compared with the upper bits of the address. For a 32-bit address, this comparison encompasses bit 31 through bit $m+b$. If the comparison reveals that the contents are identical, the comparison is considered to be successful. In common parlance, this is referred to as a "hit" in the cache. Accordingly, the value read from the data array may be used. Otherwise, if the comparison does not reveal a match, the comparison is not considered to be successful. In this case, it is said that there has been a "miss". As a result, the processor must take further action, which typically includes reloading the cache entry.

[0022] As should be apparent to those skilled in the art, and for purposes of the discussion herein, a “hit” refers to finding a valid entry within a cache with a matching tag. A “miss” refers to an entry that is either not valid within the cache or whose tag does not match. Where the entry is not valid, further processing is required, including a reloading of the cache entry, as noted above.

[0023] If an entry in a cache is not valid or does not have a matching tag, the cache is reloaded. Reloading an entry involves: (1) reading 2^b bytes from the memory location specified by the address, (2) writing that data (i.e., the 2^b bytes) into the data array at the location specified by bits $m+b-1$ to b of the address, and (3) writing the upper bits of the address into the tag array.

[0024] As should be appreciated by those skilled in the art, given the tag value for any specific entry, it is possible to determine the memory address to which the entry corresponds. In other words, it is possible to determine the value from a specific memory location that is being stored in the data array. In addition, as should be apparent, for multiple entries, it is possible to determine the values from several memory locations that are being stored in the data array.

[0025] When the processor is started, or after it is reset, the tag and data arrays typically hold random values. In particular, the data value of any entry may not match the contents of the memory address determined by the tag value for that entry. This mismatch presents a problem, which is addressed by use of the valid array.

[0026] The valid array includes a bit for each entry in the cache. On reset or on processor start-up, all of the bits in the valid array are cleared to 0. If a cache entry is accessed, and the bit in the valid array is 0, access of the cache entry is treated as a “miss”. As noted above, the cache entry is then reloaded. After the cache entry is reloaded, the valid bit is set to 1.

[0027] Once a valid bit is set to 1, that bit is not set to 0 or reset to 0 until the next reset instruction or start-up instruction is executed. Thus, the valid array provides a useful function only until every cache entry has been reloaded once, if necessary. After each of the cache entries have been reloaded one time, all of the entries in the valid array will be set to 1. Once the valid array reaches this state or condition, any further read of the valid array serves no further use. In other words, once all of the bits in the valid array

are set to 1, there is no longer any reason to read the valid array. All subsequent reads will return the same result.

[0028] As noted above, in the prior art, it is customary for the valid array to be accessed even after all of the bits in the valid array have been set to 1. As should be apparent, these reads of the valid array consume power needlessly.

[0029] The invention avoids needless reading of the bits in the valid array by providing, among other features, a valid control bit. The valid control bit for each cache may be stored in a control register. The program being executed writes to the control register during execution. On startup or after a reset instruction is executed, the valid control bit is set to 0.

[0030] When the valid control bit is 0, the program recognizes that all of the entries in the cache have not been validated. As a result, the program operates according to three control parameters. First, the valid array is read for every cache access. Second, if a bit in the valid array is set to 0, the access is treated as a miss. Third, if the bit in the valid array is set to 1, the access is treated as a hit or a miss based on the tag comparison. In other words, before the valid control bit is set to 1, the program operates according to traditional guidelines, at least with respect to the valid array.

[0031] When the valid control bit is set to 1, however, the program operates according to a different set of behaviors. First, the valid cache array reads are disabled. Second, the access is treated as a hit or a miss based on the tag comparison.

[0032] One possible implementation of the hit or miss logic in pseudo-code of the cache access using the control bit is shown below in Code Segment #1. For purposes of Code Segment #1, it is assumed that the addresses are 32 bits long and that and the cache includes 256 entries of 16 bytes each.

Code Segment #1

```
entry = address[11:4];
if( valid_control = 1 )
    valid_bit = 1;
else
    valid_bit = valid_array[entry];
if( address[31:12] == tag_array[entry] )
    tag_hit = 1;
```

```
else
  tag_hit = 0;
hit = tag_hit & valid_bit;
```

[0033] To make use of the valid control bit, the processor must ensure that all the bits in the valid array are 1. The processor may do this by executing a program that touches or examines all of the entries in the cache. For an instruction cache, this involves running a code sequence that executes an instruction from every entry in the cache. For a data cache, this involves running a code sequence that reads a memory location from every entry in the cache. After the program has run either of these code sequences, it may set the valid control bit for the corresponding cache.

[0034] In an alternative to the architecture and method described above, it is contemplated that the valid control bit and the valid array entries may operate such that the use of 0s and 1s is reversed. In this alternative approach, the valid control bit initially may be 1. In this instance, once all of the entries are validated, the valid control bit may be set to 0. For the valid array, all of the entries may be initialized to zero. In this case, once an entry is validated, the entry may be changed to 0. Other variations are also contemplated to fall within the scope of the invention.

Set Associative Caches

[0035] Set associative caches present additional challenges with respect to the method of the invention.

[0036] A set associative cache may be treated quite similarly to the direct mapped cache described above. In an N-way set associative cache, there are N sets of data arrays with tag arrays and valid bits for each entry. Typically, all of these arrays are accessed in parallel, and each of the N tag values are compared with upper bits of the address. The N valid bits are used to determine if the corresponding entry is valid. Generally, at most one of the valid tag comparisons is expected to succeed in this scenario. If so, it is treated as a hit and the data array value corresponding to the matching tag is used. Otherwise, the cache access is a miss, and the address is reloaded into the cache. One set of tag arrays or data arrays for that entry is selected and the reloaded address or data is written to that set.

[0037] For our solution to be applicable to a set associative cache, it is desirable to establish a sequence that guarantees that all entries in the cache have been reloaded. Knowledge of the replacement policy (i.e., when reloading, which set will the reload go into) is key to generating such a sequence. For instance, if the reload policy is completely random, there is no sequence that guarantees that all sets in all entries will be accessed, and thereby set to a valid condition or state.

[0038] For a common replacement policy, known as LRU (least recently used), the sequence for setting all bits to a valid condition is straight forward. For a N-way set associative cache, the program accesses N addresses corresponding to each entry.

Flush

[0039] Some processors support a cache flush instruction, which invalidates a cache entry. If this instruction is implemented by means of clearing the valid bit for that cache entry, the program additionally ensures that the valid-control bit is set to 0. It may then run an additional code sequence to ensure that the invalidated bit is reset to 1, and then set the valid-control bit back to 0.

Hardware Controlled

[0040] While the focus of this discussion has been on the concept of using a software based valid control bit, the same results may be achieved using hardware. As should be appreciated by those skilled in the art in view of the foregoing discussion, hardware may be designed where the valid control array ANDs all of its bits together. If the result of the AND operation is 1, then all bits in the array are 1, and it is no longer necessary to access the array. As may be appreciated by those skilled in the art, an OR operand may be used in the place of the AND operand without departing from the scope of the invention.

[0041] Reference is now made to Figs. 1-3 of the drawings. Figs. 1-3 illustrate one embodiment of the invention.

[0042] With reference to Fig. 1, the method 10 begins at 12. The method 10 is intended to be executed by a processor for accessing a cache. As noted above, a cache includes a tag array, a data array, and a valid array. At 14, the method 10 sets all valid

array entries in the valid array to a first array value on processor reset or startup. The method 10 proceeds to 16, where it reads a valid array entry from the valid array each time a data array entry from the data array is accessed. After 16, the method 10 selects between two separate paths. At 18, if the valid array entry is a first array value, access to the cache is treated as a miss, which indicates that the entry is invalid. If so, at 20, the cache is reloaded from memory and the valid array entry is set to the second array value. Following the other path, at 22, if the valid array entry is a second array value, a tag array entry in the tag array is compared with selected bits in an address to determine if the data array entry is a hit or a miss. A hit indicates a valid entry while a miss indicates an invalid entry. In Fig. 1, a transition block 24, also labeled "A", is provided to indicate that the method 10 continues in Fig. 2.

[0043] In Fig. 1, the dotted border 26 that encompasses 16, 18, 20, and 22 is intended to convey that these four operations are performed during an interval before all of the entries in the valid array have been validated.

[0044] With reference to Fig. 2, the method 10 proceeds to 28 where a valid control register is established for the cache. At 30, the valid control register is set to a first control value before all valid array entries are validated. At 32, the valid control register is set to a second control value after all valid array entries are validated. A second transition block 34, also labeled "B", indicates that the method 10 continues in Fig. 3.

[0045] With reference to Fig. 3, the method 10 continues at 36 where, after the second control value is set, reads of the valid array are disabled. At 38, after reads of the valid array are disabled, the tag array entry is compared with the selected bits in the address to determine if the data array entry is a hit or a miss. In the case of a miss, the cache is reloaded from memory at 40. The method ends at 42. The dotted border 44 is intended to indicate that the operations 36, 38, and 40 are executed during an interval after all of the valid array entries have been validated.

[0046] As noted above, in one contemplated variation of the method 10, the data array may comprise 2^m entries, each of the data array entries may comprise 2^b bytes, and there may be a total number of addresses, 2^a . Of course the data array and the entries also may be of a larger or smaller size without departing from the scope of the invention.

[0047] In a variation of the method 10, the tag array is compared with upper bits in the address. Alternatively, where 2^m entries with 2^b bytes are used, it is contemplated that the tag array entry is compared with bits a-1 to m+b of the address. If the address is a 32-bit address, the tag array is compared with bits 31 to m+b in the address.

[0048] With respect to reloading of the data entries, reloading is contemplated to involve several operations. First, 2^b bytes are read from a memory location specified by the address. Second, the 2^b bytes are written to a data array entry specified by bits m+b-1 to b of the address. Third, upper bits, a to m+b, from the address are written to the corresponding tag array entry. Other contemplated reload operations also are contemplated to fall within the scope of the invention.

[0049] With respect to the valid control register, it is contemplated that the register is a one bit register. Similarly, it is contemplated that all of the valid array entries are single bit entries.

[0050] For purposes of the invention, the cache may be a data cache, or it may be an instruction cache. The cache also may be a set associative cache with N sets of data arrays. If so, each of the N sets of data arrays may be accessed in parallel.

[0051] With respect to the method 10, if, after the valid array entry has been set to the second array value, the data array entry subsequently becomes invalid, the method 10 contemplates resetting the valid array entry to the first array value. As should be appreciated from the foregoing, if the valid control register contains the second control value, the valid control register is reset to contain the first control value.

[0052] In a variation contemplated for the method 10, the second control value is established by combining all of the valid array entries.

[0053] As may be appreciated from the discussion above, the method 10 contemplates that the all of the valid array entries are combined via hardware. Where hardware is employed, all of the valid array entries may be combined via an AND operand or an OR operand. Where an AND operand is employed, the first array value is 0, the second array value is 1, the first control value is 0 and the second control value is 1. Where an OR operand is used, the first array value is 1, the second array value is 0, the first control value is 1 and the second control value is 0.

[0054] In another contemplated variation of the method, the second control value is established in response to an instruction executed by the processor.

[0055] The method 10 also may include a validation forcing operation. A validation forcing operation forces validation of the valid array. A validation forcing instruction includes reading a predetermined sequence of addresses, setting all of the valid array entries to the second array value, and setting the valid control register to the second control value. In a variation of the validation forcing operation, instructions may be executed in a predetermined sequence of addresses, thereby setting all valid array entries to the second array value. Then, the valid control register may be set to the second control value.

[0056] Other embodiments and variations of the invention, whether via a method or via hardware, should be apparent to those skilled in the art based on the embodiments described above. Those variations and equivalents are intended to fall within the scope of the invention.

What is claimed is:

1. A method executed by a processor for accessing a cache, wherein the cache comprises a tag array, a data array and a valid array, the method comprising:
 - setting all valid array entries in the valid array to a first array value on processor reset or startup;
 - before all valid array entries are validated,
 - reading a valid array entry from the valid array each time a data array entry from the data array is accessed,
 - if the valid array entry is the first array value, treating access to the cache as a miss
 - if the valid array entry is a second array value, comparing a tag array entry in the tag array with selected bits in an address to determine if the data array entry is a hit or a miss, and
 - in case of a miss, reloading the cache from memory, and setting the valid array entry to the second array value;
 - establishing a valid control register for the cache;
 - setting the valid control register to a first control value before all valid array entries are validated;
 - setting the valid control register to a second control value after all valid array entries are validated; and
 - after the second control value is set,
 - disabling reads of the valid array
 - comparing the tag array entry with the selected bits in the address to determine if the data array entry is a hit or a miss, and
 - in case of a miss, reloading the cache from memory
2. The method of claim 1, wherein the data array comprises 2^m entries, each of the data array entries comprise 2^b bytes, and the total number of addresses is 2^a , where m is an integer, a is an integer and b is an integer.

3. The method of claim 1, wherein the valid control register comprises one bit.
4. The method of claim 1, wherein all of the valid array entries comprise one bit.
5. The method of claim 2, wherein the tag array entry is compared with bits a-1 to m+b of the address.
6. The method of claim 5, wherein:
the address is a 32 bit address, and
the tag array is compared with bits 31 to m+b in the address.
7. The method of claim 1, wherein the tag array is compared with upper bits in the address.
8. The method of claim 1, wherein the reloading comprises:
reading 2^b bytes from a memory location specified by the address;
writing the 2^b bytes to a data array entry specified by bits m+b-1 to b of the address; and
writing upper bits a to m+b from the address to the corresponding tag array entry,
wherein a, b and m are all integers.
9. The method of claim 1, wherein the cache is a data cache.
10. The method of claim 1, wherein the cache is an instruction cache and the data array comprises instructions.
11. The method of claim 1, wherein the cache is a set associative cache with N sets of data arrays, with N being an integer, and wherein each of the N sets of data arrays are accessed in parallel.
12. The method of claim 1, further comprising:

if, after the valid array entry has been set to the second array value, the data array entry subsequently becomes invalid,

resetting the valid array entry to the first array value; and

if the valid control register contains the second control value, resetting the valid control register to contain the first control value.

13. The method of claim 1, wherein the second control value is established by combining all of the valid array entries.

14. The method of claim 1, wherein the all of the valid array entries are combined via hardware.

15. The method of claim 14, wherein the first array value is 0, the second array value is 1, the first control value is 0, the second control value is 1, and all of the valid array entries are combined via an AND operand.

16. The method of claim 14, wherein the first array value is 1, the second array value is 0, the first control value is 1, the second control value is 0, and all of the valid array entries are combined via an OR operand.

17. The method of claim 1, wherein the second control value is established in response to an instruction executed by the processor.

18. The method of claim 17, further comprising:
executing a validation forcing instruction comprising
reading a predetermined sequence of addresses,
setting all of the valid array entries to the second array value, and
setting the valid control register to the second control value.

19. The method of claim 17, further comprising:
executing a validation forcing instruction comprising

executing instructions in a predetermined sequence of addresses, thereby setting all valid array entries to the second array value, and setting the valid control register to the second control value.

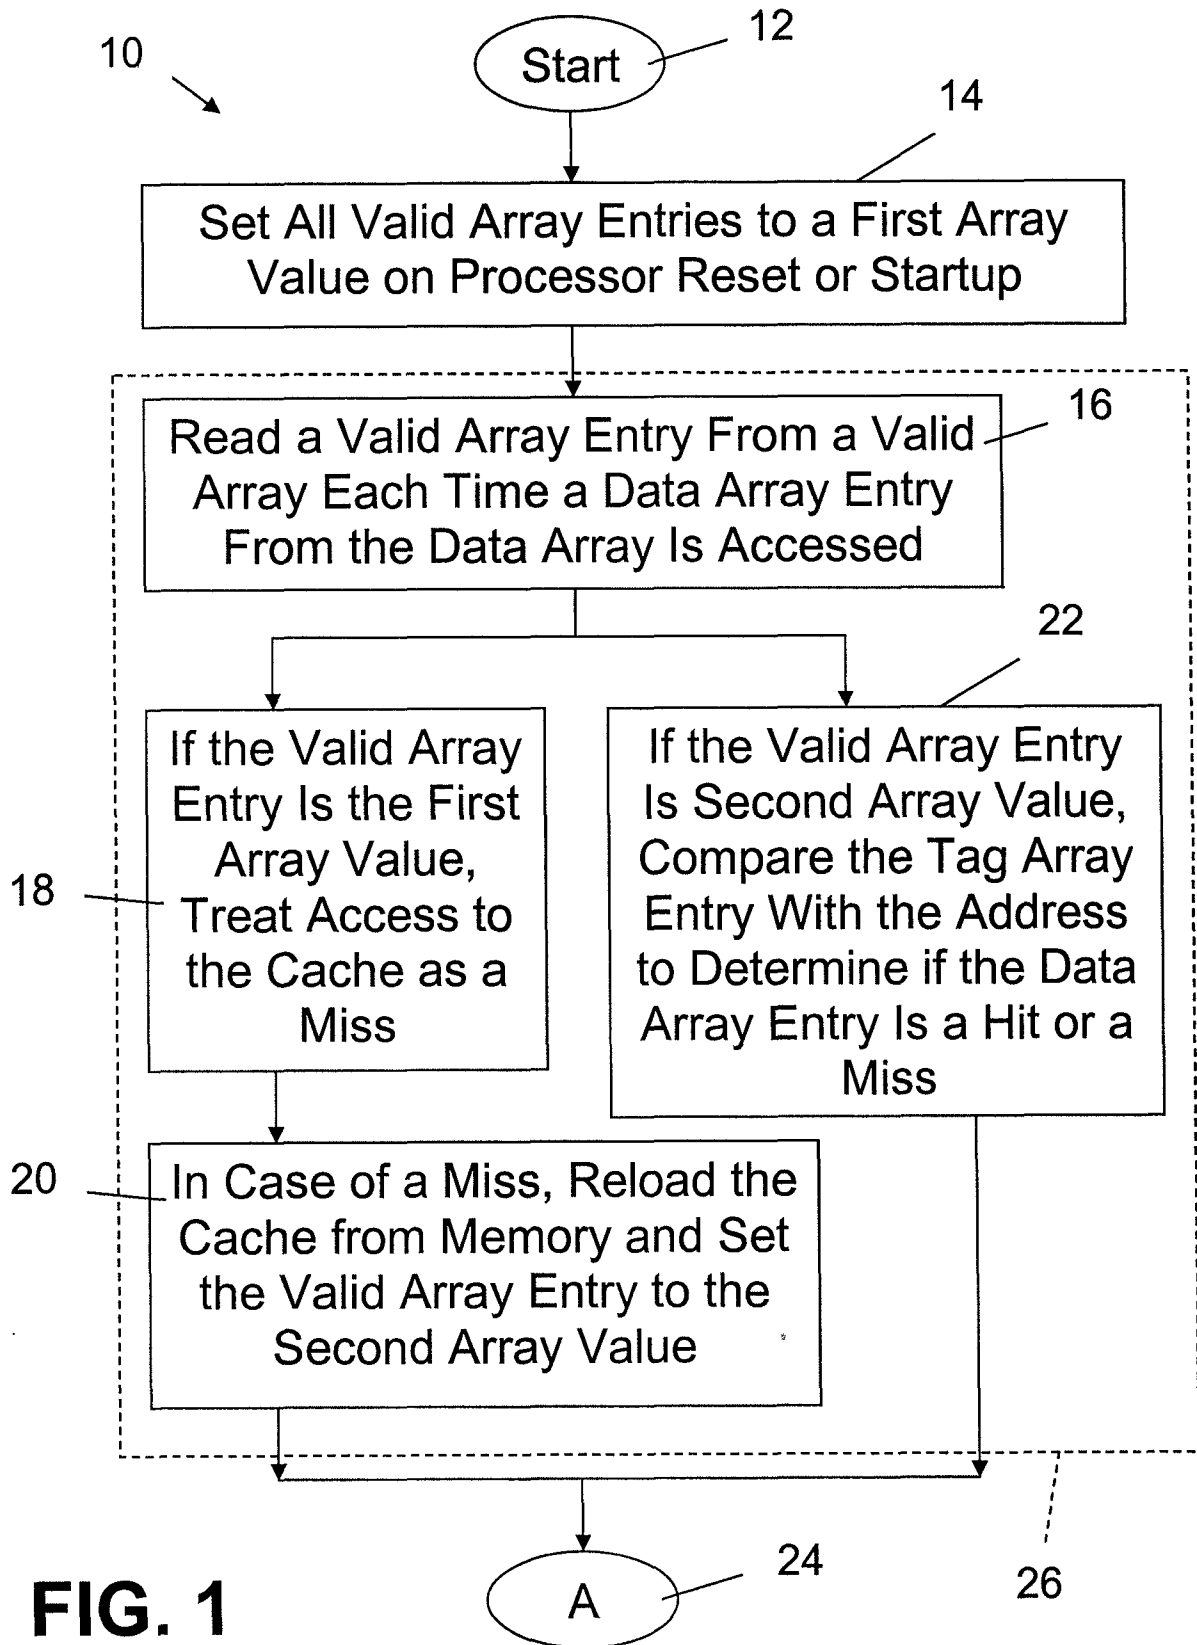


FIG. 1

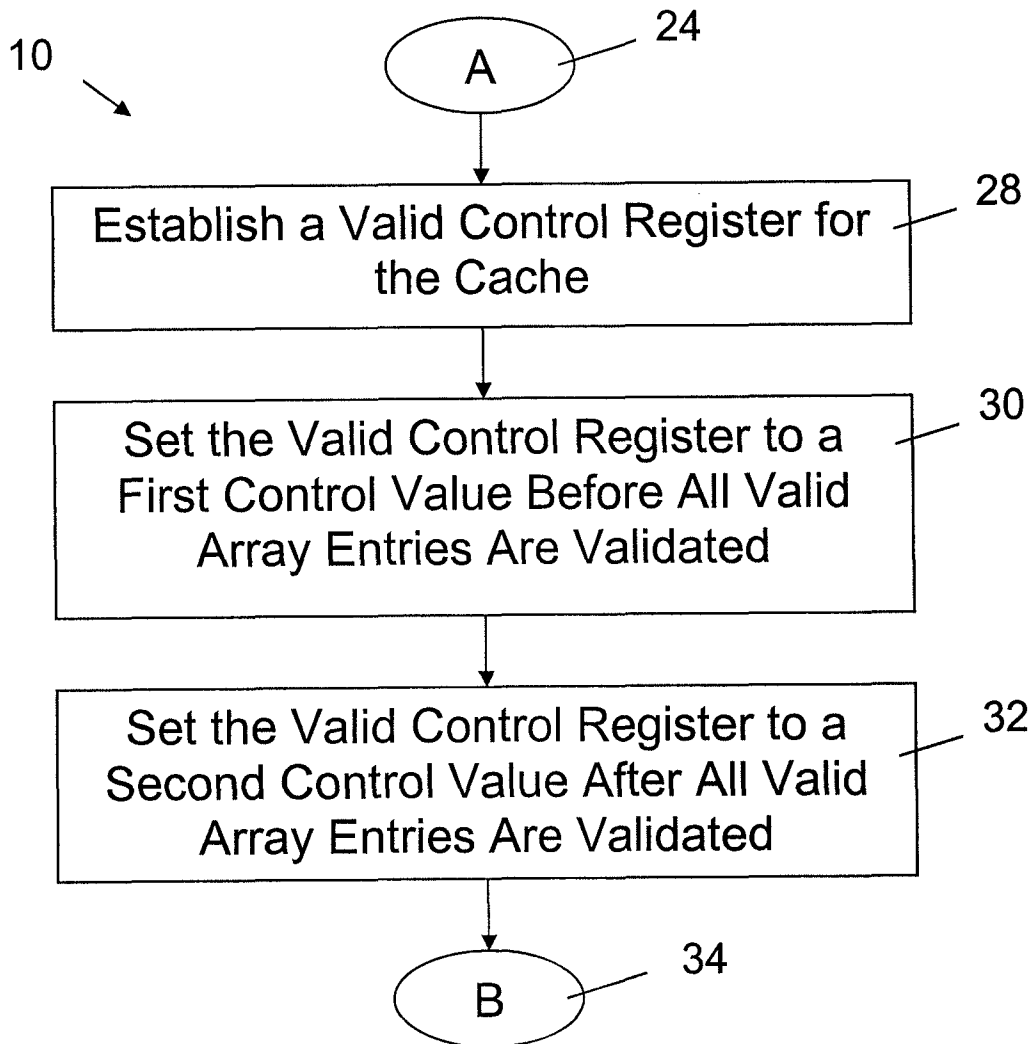


FIG. 2

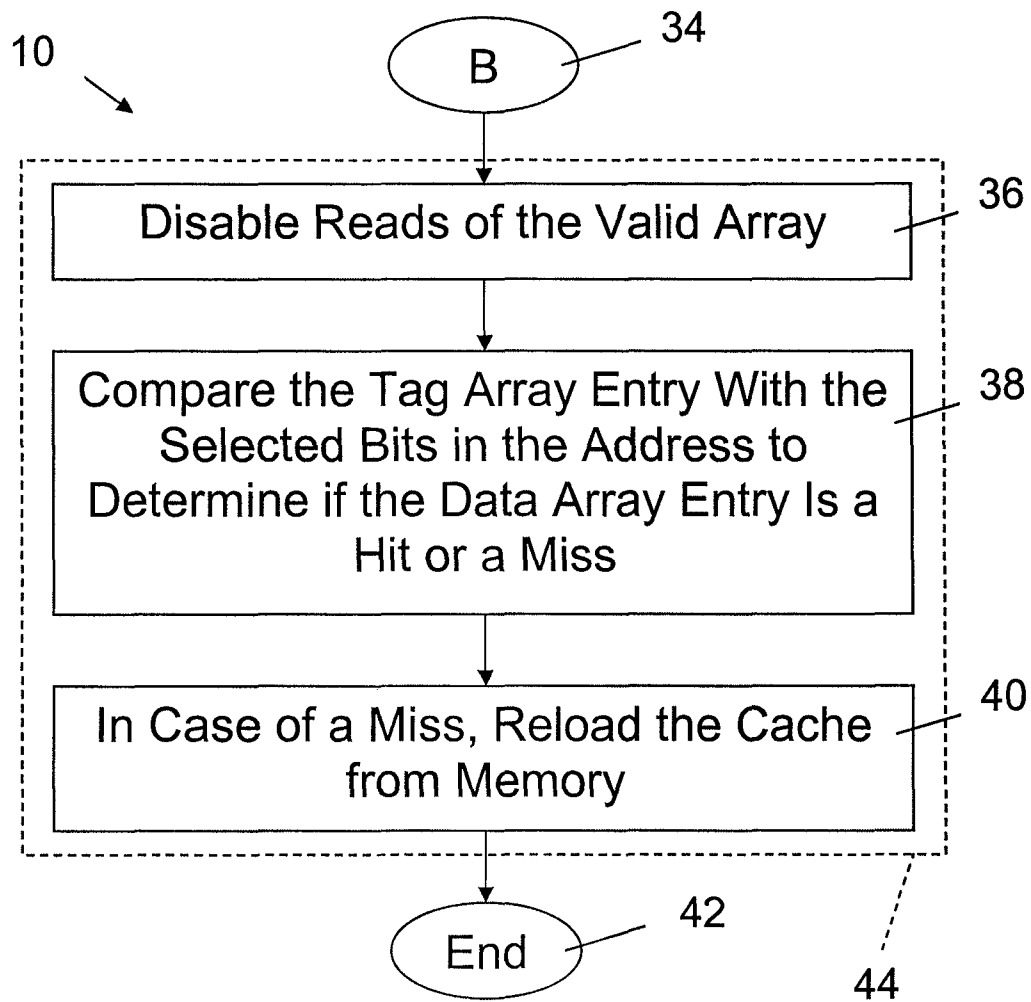


FIG. 3