



(51) International Patent Classification:

G11C 15/04 (2006.01) G11C 7/10 (2006.01)
G11C 16/04 (2006.01) G06F 15/78 (2006.01)

(21) International Application Number:

PCT/US2013/068420

(22) International Filing Date:

5 November 2013 (05.11.2013)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/724,401 9 November 2012 (09.11.2012) US
61/730,884 28 November 2012 (28.11.2012) US
13/749,361 24 January 2013 (24.01.2013) US
13/756,076 31 January 2013 (31.01.2013) US
13/756,131 31 January 2013 (31.01.2013) US

(71) Applicant: **SANDISK TECHNOLOGIES INC.**
[US/US]; Two Legacy Town Center, 6900 North Dallas
Parkway, Plano, Texas 75024 (US).

(72) Inventors: **SPROUSE, Steven T.**; 5285 Hoyet Drive, San
Jose, California 95129 (US). **LI, Yan**; 695 Kevenaire
Drive, Milpitas, California 95035 (US). **GEORGE, Jo-**
hann; 710 Russett Terrace, Sunnyvale, California 94087
(US).


(74) Agents: **CLEVELAND, Michael G.** et al.; Davis Wright
Tremaine LLP, 505 Montgomery Street, Suite 800, San
Francisco, California 94111 (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, BN, BR, BW, BY,
BZ, CA, CH, CL, 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, IR, IS, JP, KE, KG, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, 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, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: ON-DEVICE DATA ANALYTICS USING NAND FLASH BASED INTELLIGENT MEMORY

**FIG. 23**

(57) Abstract: A NAND Flash based content addressable memory (CAM) is used for a key-value addressed storage drive. The device can use a standard transport protocol such as PCI-E, SAS, SATA, eMMC, SCSI, and so on. A host writes a key-value pair to the drive, where the drive writes the keys along bit lines of a CAM NAND portion of the drive and stores the value in the drive. The drive then maintains a table linking the keys to location of the value. In a read process, the host provides a key to drive, which then broadcasts down the word lines of blocks storing the keys. Based on any matching bit lines, the tables can then be used to retrieve and supply the corresponding data to the host. The system can be applied to perform a wide range of analytics on data sets loaded into the NAND array.



Published:

- *without international search report and to be republished
upon receipt of that report (Rule 48.2(g))*

ON-DEVICE DATA ANALYTICS USING NAND FLASH BASED INTELLIGENT MEMORY

FIELD OF THE INVENTION

[0001] This invention relates generally to non-volatile memories and, more specifically, to using non-volatile memory of a NAND-type architecture perform on-chip data analytics.

BACKGROUND OF THE INVENTION

[0002] Content addressable memories, also known as associative memories, are different from standard memories in the way that data is addressed and retrieved. In a conventional memory, an address is supplied and the data located at this specified address is retrieved. In contrast, in a content addressable memory (CAM), data is written as a key-data pair. To retrieve the data, a search key is supplied and all the keys in the memory are searched for a match. If a match is found, the corresponding data is retrieved.

[0003] Content Addressable Memories, or CAMs, can be implemented in several ways. In one sort of embodiment, a CAM is implemented using a conventional memory and an associated CPU which searches through the memory to find a matching key. The keys in the memory may be sorted, in which case a binary search can be used; or they can be unsorted, in which case they are usually hashed into buckets and each bucket is searched linearly. A CAM can also be implemented as a semiconductor memory, where every memory location contains an n-bit comparator. When an n-bit key is provided, each entry in the CAM will compare the search key with the entry's key, and signal a match if the two are equal.

SUMMARY OF INVENTION

[0004] A first set of aspects concern a method of analyzing data. A plurality of data sets selected by a host are written in a memory array of a NAND architecture of NAND strings formed along bit lines and spanned by word lines, where the data sets are written oriented along the bit lines with corresponding data elements arranged to line along the same word lines. A plurality of word lines are concurrently biased to a first set of

sensing voltage levels corresponding to host determined parameter values. The method then determines those of the bit lines that conduct in response to said biasing of a plurality of word lines to the first set of sensing voltage levels.

[0005] Other aspects relate to a method of operating a non-volatile array of a NAND architecture, the array having multiple of blocks of NAND strings formed along bit lines and spanned by word lines, where the bit lines span a plurality of the blocks and are each connected to a corresponding set of latches. The method includes reading a first page of data from a first word line of the array and storing the first page of data in a first of the latches for each of the bit lines of the array. A second page of data can be read from a second word line of the array and stored in a second of the latches for each of the bit lines of the array. Within the sets of latches, one or more arithmetical/logical operations upon the contents of the first and second latches are preformed, with the result subsequently being written from the sets of latches along a third word line of the array.

[0006] Various aspects, advantages, features and embodiments of the present invention are included in the following description of exemplary examples thereof, which description should be taken in conjunction with the accompanying drawings. All patents, patent applications, articles, other publications, documents and things referenced herein are hereby incorporated herein by this reference in their entirety for all purposes. To the extent of any inconsistency or conflict in the definition or use of terms between any of the incorporated publications, documents or things and the present application, those of the present application shall prevail.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Figure 1 is a schematic representation of a NAND array used as a CAM memory.

[0008] Figure 2 is a schematic illustration of the network of some of the elements to supply the word line in a NAND array for conventional operation.

[0009] Figure 3 is a schematic illustration of the network of some of the elements to supply the word line in a NAND array for CAM operation.

[0010] Figure 4 shows one embodiment for how keys can be written along bit lines of

an NAND array and searched.

[0011] Figure 5 given some detail on how a key/inverse pair from Figure 4 is programmed into a pair of NAND strings.

[0012] Figures 6A-C shows another embodiment for how keys can be written along bit lines of an NAND array and searched.

[0013] Figure 7 shows an exemplary encoding of 2-bits per cells for four state memory cell operation.

[0014] Figure 8 shows how the data states and the complementary data used for the inverted keys correspond in the 2-bit per cell example.

[0015] Figure 9 shows an example of how a key would be encoded onto a 4 cell NAND string on bit line BL and its inverse on bit line BLB.

[0016] Figure 10 illustrates the process of matching of content in word line direction.

[0017] Figure 11 illustrates how the position of a conducting bit line can be used as an index in to another table that can be used to retrieve data associated with the target key.

[0018] Figure 12 schematically illustrates how a key-value pair is stored in a NAND based CAM and how the value is accessed using the key.

[0019] Figure 13 illustrates a memory arrangement for transposing the data keys.

[0020] Figure 14 represents a first hardware embodiment for transposing data using a FIFO-type structure.

[0021] Figure 15 represents another hardware embodiment for transposing data.

[0022] Figure 16 shows one embodiment of a memory system incorporating a CAM type NAND into a solid state drive (SSD) for performing data analytic within the memory system.

[0023] Figure 17 illustrates how data analytics with numerical range detection can be performed by exploiting an array's NAND structure.

[0024] Figure 18 is an example of data latch assignments for the process illustrated by Figure 17.

[0025] Figures 19 and 20 illustrate some steps of two search processes.

[0026] Figures 21 and 22 illustrate a maximum and a minimum search operation.

[0027] Figures 23 and 24 respectively give a schematic representation of an on-chip arithmetical operation and a corresponding latch utilization.

[0028] Figures 25A-C illustrate some detail of how arithmetic operations can be performed.

[0029] Figures 26A and 26B show how more latches can be used to perform arithmetic operations involving more n

[0030] Figures 27 and 28 illustrate an application to financial data analysis.

[0031] Figures 29-31 show some examples of how a data set can be placed on more than one NAND string and corresponding latch structures.

[0032] Figures 32 and 33 respectively illustrate digital and analog counting techniques for analytics results.

[0033] Figure 34 gives an example of file mapping for performing analytics on large file systems.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Content Addressable Memory based on NAND Flash Memory

[0034] The following presents a method of using a Flash based NAND memory array as a content addressable memory (CAM) that can be realized in both binary and ternary embodiments. As described in more detail below, keys can be programmed along the bit lines of a block. The search key is then input along the word lines of the blocks, so that a bit line on which a corresponding key has been programmed will be conducting. This allows for all the keys of a block to be checked at the same time.

[0035] The typical way by which a NAND memory array is read is that data is read out a single word line (or portion of a word line) at a time, with the non-selected word lines

along the NAND strings being biased so that they are fully turned on regardless of the data state, removing the non-selected memory from affecting the read operation. In this way, the data content of the memory is read out a page (the unit of read) at a time. In contrast, to use a NAND flash memory as a content addressable memory, all of the word lines are set to a specific data dependent value, where the data is the key, and the memory determines which bit lines then conduct, thereby determining particular bit lines correspond to the input key, rather than the data of individual cells. An operation where sensing voltages are applied to multiple word lines in the context of an enhanced post-write read operation is given in US patent application number 13/332,780 filed on December 21, 2011, (and which also presents more detail on NAND flash memory in general); however, even in that case only a few of the word lines receive a sensing voltage. Also, in prior art NAND memories, data was aligned along word lines, where data pages (for both read and write) are aligned along the word lines. Here, data is aligned along bit lines and many, or even all, of the word lines along the bit lines can receive either a high voltage sufficient to turn on a cell in a programmed state, or a low voltage sufficient to turn on a cell in the erased state. The following discussion will use the EEPROM based flash memory as the exemplary embodiment, but other memory devices having a NAND type of architecture, including 3D NAND (such as described in T. Maeda et al., "Multi-stacked 1G cell/layer Pipe-shaped BiCS flash memory", 2009 Symposium on VLSI Circuits, pages 22-23) for example, can also be used.

[0036] In a binary, EEPROM based flash memory, in a write operation each cell is either left in an erased state or charge is placed on the cell's floating gate to put the cell in a programmed state, which here are respectively taken as the 1 and 0 states. When a low value for the read voltage is applied to its control gate, only a cell in the erased, or 1, state will conduct. For cells in the programmed, or 0, state, a high value of the read voltage needs to be applied to the control gate for a cell to conduct. The keys will be arranged along bit lines of a block of the memory array. Since a cell in the 1 state will conduct for either read voltage, each key needs to be written twice, in inverted and non-inverted form. As discussed below, this can be done by writing the target key along one bit line and its inverse along another, or writing half the bit line with the (non-inverted) target key and the other half of the bit line with the inverted target key. More key info can be compressed into the NAND chain using multiple bits programming. For example, in a 2-3 bits per cell case, the key can be sorted in the

controller RAM and the bits will be programmed as lower, (middle) or upper pages. The following discussion will mostly be given in terms of a binary embodiment, with some specifics of the multi-state case are discussed later.

[0037] The general concept can be illustrated by Figure 1. Target keys Key 0, Key 1, ... are programmed down bit lines BL0, BL1, ... of a NAND block. Data is programmed in a separate location that can be indexed by the target key's column address number. To search the block for a key, the search key is broadcasted on the block's word lines by setting all of the word lines according to either the high or low read voltage according to the search key. (In addition to setting the word line voltages according to the key, the select gates at the end of the NAND string will also need to be turned on.) Each BL effectively compares itself to the WL key pattern for all of the bit lines in the block at the same time. If the bit line key matches the search key, the whole of the bit line will be conducting and a "1" will be read out. (Note that, as discussed further, this discussion is somewhat simplified for the reasons discussed in the last paragraph.) Once the column index of the key is found, it can be used to fetch the corresponding data from a "data" block. The key can be the hash code of the data page that will lead to the right data page by the column address of the matched NAND chain. For content matching applications, such as data compression or de-duplication, each 16KB, say, of content can generate a corresponding hash code that can be stored along the NAND chain. If the key along the NAND chain is matched, then the data page will be compared with the comparing data along the word line to avoid hash collision cases. In other cases, the content along the word line may not be a hash value, but characteristics of the data elements that can be searched as a keys to data; or the bits lines themselves main be the elements of the data themselves, rather than a pointer to a data base.

[0038] Under the arrangement illustrated by Figure 1, all of the bit lines of the array, and consequently all of the keys, are searched at the same time. In arrays that do not use an all bit line type of architecture, the number of keys searched simultaneously would be the number of bit line sensed in parallel, such as half of the total in an odd-even arrangement. The size of the key is the number of word lines. In practice, these maximum values of the keys will typically be somewhat less, since some column are usually set aside for defects, for instance.

[0039] As noted above, since a memory cell in either the 0 or 1 state will conduct for a high read voltage, the key will need to be entered twice, both non-inverted and inverted. This can be done by either programming the target key on two bit lines, reducing the number of keys by half, or programming both versions of the key on the same bit line, reducing the key size by half. However, given the size of available NAND blocks, even with these reductions the number of keys that can be checked in parallel is quite large. Relative to some other memory technologies, NAND flash memory has relatively large latencies in its operation, but in many applications this would more than be offset by the number of keys (bit lines) that can be checked in parallel (128K, for example). The process can all be done on-chip and, as only the bit lines that meet the matching case conducting current, with relatively low power consumption, so that compared to toggling out all of the data from the memory and doing the compare in the controller, it is a process of relatively low power and higher speed.

[0040] Looking at some implementation detail, an exemplary embodiment can be based on a flash memory where the indices are saved on the 128Gb NAND chains. An all bit line (ABL) architecture is used where one sensing operations will perform a match operation on all of the indices on a block at the same time. Extra column redundancy is included to avoid any bad columns (more detail on such redundancy and the accessing of columns, as well as flash memory in general, can be found in the following US patent publication/application numbers: US-2005-0141387-A1; US-2008-0266957-A1; US-2011-0002169-A1; US-2010-0329007-A1; 13/463,422; and 13/420,961.) Two copies of the same data, Data and Data Bar, are written into the NAND chain. In the example, this allows for $16\text{KB}/2/2=32000$ sets of information with a 128 bit key.

[0041] When writing in the keys, these will be typically written on a page by page basis, although in memories that allow it, partial page programming can be used to write part of the keys, with more added later. Such partial page programming is typically more limited for multi-states implementations than in binary blocks. As one example, the data can be shifted on to the memory and the inverted data can be generated on the memory to save effort on the controller for these data manipulations, where the data and data bar can be written without shifting in the data twice, with the data being written first, and the generated inverse next. Both the keys and the data can

be input into the memory system, or in some cases the keys could be generated on the memory system by the controller from the data, such as by generating hash values from the data to use as keys. If the keys are to be sorted before being written along the bit lines, this will typically be done on the controller due to the amount of data involved, such as multiple blocks' worth of data. For example, the data could initially be written in a particular area, say die 0, plane 0, blocks 0-15, and then sorted and written into the blocks having been sorted to the block level. Alternately, the keys could be assembled in RAM (either on the controller or on a separate chip) or cache NAND memory (such as described in US provisional application number 61/713,038) before sorting them to the desired level of granularity and writing them into a set of blocks.

[0042] As discussed further below, the data/data bar pairs can be written on two bits lines or on a single bit line. When the data/data bar pairs are written on two bit lines, such as discussed with respect to Figure 4, the pairs can be written next to each other or in other patterns, such as writing the data bit lines in one area and the inverted data bit lines in another zone. When both parts of the pair are written on the same bit line, as discussed below with respect to Figure 6A, they can be written in a top/bottom format or interleaved. For example, when the data and inverted data are interleaved to alternate down the word lines, this has the advantage that at most two elements in a row are the same down the bit line; further, interleaving can lead to efficient data transfer on to the memory as first a page of data is transferred on the memory and the next page can just be generated in the latches by inverting all the bits, as the next page is the inverted data of the first page.

[0043] The matched index can then be linked to other data corresponding to the determined column address; for instance, the keys could be a hash value, such as from a Secure Hash Algorithm (SHA), used to point to the actual data that can also be stored elsewhere on the memory itself. All the matching can be done inside of the NAND chip and, when the match is found, the column address can also be transferred out if needed or just the data, if also stored on the NAND chip, can be transferred out.

[0044] To efficiently implement the use of a NAND array as a CAM memory, changes can be made to the word line driving circuitry. To broadcast a search key down the word lines of a block, in addition to turning on the select gates on either end of the NAND strings, each word line of the block needs to be set to either the high or low read

voltage according to the search key. This is in contrast to typical NAND operation, where only a single word line at a time is selected for a read voltage, with all of the other word lines receiving a pass voltage sufficient to remove them from influencing the sensing regardless of their data state.

[0045] Figure 2 is a schematic illustration of the network of some of the elements to supply the word line in a NAND array for conventional operation. At 201 is the cell array for a plane of a NAND chip, with two blocks explicitly marked out at 203 and 205. Each block's word lines are feed by a word line select gate WLSW 213 or 215 as controlled from select circuitry at 217. The bit lines are not indicated, but would run down to the sense amp block S/A 207. The various control gate voltage CGI are then supplied to the select gates 213 and 215 from the drivers CG drivers 231 and UCG drivers 233 and 235 by way of switches 223 and 225, respectively. In the exemplary embodiment shown here, a block is taken to have 132 word lines, where a pair of dummy word lines are included on both the drain and source sides of the NAND strings. The UCG Drivers 233 and 235 are for supplying the pass voltages used on unselected word lines during program, (standard, non-CAM) read or verify operations. As this level is used on the large majority of word lines, these can be lumped together for a single driver. The selected control gates are biased to VPGM at program, CGR voltage at read or verify. In Figure 2, CGI<126:1> is the decoded global CG lines. CGI<0> and CGI<127>, that are here biased differently from other 126 word lines due to edge word line effects. The dummy word line bias CGD0/1 is for the drain side dummy word lines and CGDS0/1 is for the source side ones.

[0046] For a typical NAND memory operation, only a few word lines at a time are individually biased. In addition to a selected word line, adjacent or edge word lines may receive special bias levels to improve operations. Consequently, existing word line drivers are arranged so that they can only take care of a handful of word lines. With logic changes, it may be possible to drive up to perhaps two dozen or so word lines. However, to drive all the word lines of a block (here 128, ignoring dummies) will require additional analog drivers. Figure 3 illustrates some of these changes.

[0047] The array 301, blocks 303 and 305, select circuitry 317, CG Drivers 331, and switches 313 and 315 can be the same as in Figure 2. The additional word line drivers are shown at 343 and 345 and can supply the word lines through respective switches at

353 and 355. In each of 343 and 345, the level shifter HVLSHIFT receives the voltage VREAD and a digital value DFF(0/1) for each word line. The level shifter then converts the digital values of 0, 1 for the broadcast key to the analog high and low word line levels. As the memory cells will still need to be written (both programmed and program verified), the other circuit sketched out in Figure 2 will still be present, though not shown in Figure 3 to simplify the discussion. It may also be preferable to make some changes to the sensing circuitry S/A 307 to more efficiently perform the XOR operation described below between the pairs of bit lines holding a key and its inverse.

[0048] Figure 4 shows the encoding of the keys along bit lines, where the key is entered twice, in non-inverted and inverted form. Here the bit lines are labeled BL for the non-inverted key and BLB for the inverted version. Here the pairs are shown as being adjacent, although this need not be the case, but will typically make XOR-ing and keeping track of data easier. Also, this arrangement readily lends itself to NAND arrays using an odd/even BL arrangement. As shown in the half of Figure 4, for reference a key of all 1s is written along BL1 and a key of all 0s is written along BLn, with the corresponding inverted keys at BLB1 and BLBn. For the defective bit lines, the bit line either stuck “0” or stuck “1” regardless of the word line voltage bias. The XOR results between the two read results will always yield “1”. The BL and BLB data pattern will eliminate the defected bit lines from yielding match results mistakenly. In this example, only seven word lines are used. A more interesting key of (1001101) is entered on BLn+1, with its inverted version at BLBn+1, as also illustrated in Figure 5.

[0049] Figure 5 shows the two corresponding NAND strings, where 0 is a programmed cell, 1 a cell left in its erased state, the cells being connected in series down the NAND strings to the common source line CELSRC. To search for this key, it is encoded as low read voltage for the 0 entries and high read voltage for the 1s. The search key is shown at the left of the top of Figure 5. When put onto the word lines, this correspondingly finds that BLn+1 is conducting (and BLBn+1 is non-conducting), as shown by the “c” (and “nc”) in the sense 1 row. However, BL1 and BLBn are also both conducting, as a cell in the 1 state will conduct for either read value.

[0050] The second sensing (these can be performed in either order) is then made with the search reversed. Although BL1 and BLBn are still conducting, the result from the key actually sought has changed: BLn+1 is now non-conducting and BLBn+1

conducts. By taking the result of the two reads and XOR-ing them, the sought key will give a 0 on the corresponding bit line and also on its inverted version. Consequently, by searching for the 00 pattern in the XOR data, the output column address can be found and the corresponding data block accessed. Under the sort of embodiment used in Figure 4, two reads are needed for the pattern match and internal pattern detection on the NAND device can judge if there is a match. The redundancy of the BL/BLB pairs provides redundancy to help protect from bad bit lines, but a second pair can also be kept for further protection. A copy of the key can also be kept with any associated data and used to check the match, where this copy can be ECC protected. Additional protection can also be provided by each bit line including several (8, for example) parity bits, for error detection and correction purposes, where the redundancy bit are preferable along the same bit lines for all of the keys so that these parity bits can either be read or taken out to the comparisons by use of a “don’t care” value applied to these word lines, as described below. For example, the data can be read when checking when checking the data, as either part of a post-write read or other data integrity check, but ignored during CAM-type operations.

[0051] Generally, for both this and other embodiments described here, a post-write read can be used to insure that the keys have been successfully written into the NAND memory, as any error bits could prevent a NAND string from conducting and would give rise to “false negatives” when matching. If an error is found, the bad data can be rewritten. In the exemplary NAND flash example, the incorrectly written data can be rewritten to another data block and any key-data correspondences updated accordingly. More detail on post-write read operations can be found in US patent application number 13/332,780 and references cited therein.

[0052] In terms of performance, in the case of a 16KB page of 128 bit keys, if two copies of the both the data and its inverse are stored, the corresponds to 4KB of keys, or 32000 keys. (As all of the word lines are sensed at once, so that here, a “page” involves a sensing of all the word lines of a block rather than a single word line.) If this page of 32000 keys is sensed in 50us, this is a rate of 0.64 GC (Giga-compares) per second per plane. If four planes are sensed in parallel, this can lead to 2.56GC/s at a consumption of about 200mW.

[0053] Figure 6A illustrates a second embodiment for how the key can be stored along

a bit line. In this case, both the key and its inverse are written onto the same bit line. For a given block, this means that the maximum key size is only half the number of word lines, but this allows for the search key and inverted key to be broadcast at the same time. Consequently, the search can be done in a single read.

[0054] Referring to Figure 6A, this shows 14 different word lines with the keys entered in the top half and the inverted versions of these same keys entered in inverted form in the bottom half of the same bit line. Thus, taking the bit line at D7, rows 1-7 hold a 7 bit key, and rows 8-14 the inverted version of the same key. (Although arranged similarly to Figure 4, in Figure 6A the top and bottom halves represent 14 different word lines where the top-bottom division is the key/inverted key boundary, whereas in Figure 4, the top and bottom are the same seven word lines repeated twice for two different sensing operations.) For comparison purposes, the keys shown in Figure 6A are the same as in Figure 4, with the bit line of D7 holding the sought for key in the top half and its inverse in the bottom half, and D8 holding the inverted key so that these two halves are switched.

[0055] To search for a key, the search pattern is then broadcast on the top half word lines and its inverse on the bottom half word lines. Any bit lines with a matching keys, in this case D7, will then conduct, as shown at bottom where “nc” is non-conducting and “c” conducting. If redundancy is desired, the non-inverted version can also be programmed in as at D8 and then detected by broadcasting the non-inverted search key, and the bit lines reads searched for a 11 pattern, which can then be output as a data pointer. If further redundancy is wanted, the key or key/inverse pair can be written into the array a second time and parity bits can also be included, much the same way as discussed for the embodiments based on Figure 4. The defective bit line should be isolated with isolation latch and not used. If some defect shows up as a stuck “0”, it can potentially generate the “false” match. In this case, the data content should be compared in order to confirm whether this is a real match or a false match. The other most common reliability issue is that some cells may have lost some charges after some time, that will also produce a “false” match. Then a content match check will eliminate the “false” match error. The word line voltage bias can be budgeted a little higher to avoid “missing” a match, which is very harmful error. A “false” match can be double checked with the content check.

[0056] Figure 6B schematically illustrates the key/inverse pairs along NAND strings. Two strings are shown (for bit lines BL_n and BL_m) each having a drain and source select gate (SGD, SGS) on either end, where the source ends are then connected along the source line CELSRC. In between are the memory cells on the strings connected in series. In this example, the strings has cell capacity to hold a 48 bit key, its 48 bit inverse, and some parity bits. Although shown here with the key along the first 48 word lines followed by the inverse along the next 48 word lines, more generally they can interleaved in various ways; for example, each of the key bits can be followed it inverse in the next word line as, when programing, this allows for a page to loading in and written, after which the programming data can be inverted in the latches and written into the next word line. The parity bits can also be variously located along the NAND string, although having them grouped can lead to easier decoding when searching the keys.

[0057] Each of bit lines BL_n and BL_m show a portion of a key along four adjacent word lines and the corresponding four adjacent word lines holding the inverse. To search the keys of the block, the word lines are then biased according to the search key, where the high sensing voltage used to checking for “0” values and the low sensing voltage to check for ‘1’ values. The high value is here taken as VREAD, and can be the same used in a typical NAND memory for non-selected word lines, and the low sensing values is labeled as V0. The select gates will also need to be on and VREAD should also be applied to the word lines holding parity bits as these as used for data integrity checks and are not meant factor into key search operations.

[0058] To make the stored keys more robust, the memory can shift the sensing margins to favor “false” matches rather than misses. (Similarly, the programming parameters can be shifter relative to those typically used.) The “false” matches can be examined by the data check later to help remove any false positives. A duplicated key can be used to check for preventing error, where these duplicates can be stored on other NAND strings, with the associated data, or other locations on the system. Relative to a standard NAND memory, this arrangement will need to add extra circuitry, as described with respect to Figures 2 and 3.

[0059] Rather than sense the search for the full key (or key/inverse) in a single sensing, a partial key can be searched, allowing the full key/inverse matching to be done

incrementally. This can allow for the less independently settable word line levels, resulting in less circuitry changes relative to a standard NAND memory, but it can require some logic changes. The full key/inverse can be searched sequentially, where each subsequent sensing will be judged based on previous sensing results. For the example of Figure 6B, rather than check all 24+24 word lines of the key/inverse in one go, a partially key check of, say 24 bits at a time can be done: if no matches are found, the process can move on to any other blocks holding keys; if a match is found, a second partial key can be checked, and so on. The subsequent checks can either do all of the NAND string again and compare the results of the partial searches, or only check those which have conducted in the previous partial key matches. Figure 6C illustrated such a partial key comparison, where only 24 bits of the 48 bits in the key are being checked. The other bits of the key and its inverse are then set to the “don’t care” value, as shown at the corresponding bits of the inverse that are set at VREAD.

[0060] As each key is written in twice (non-inverted, inverted) on a bit line, a block with 128 word lines can hold 64 bit keys, while 128 bit keys would need blocks of 256 word lines. Also, it should be noted that although the key/inverted keys are here shown as being written respectively into the top half/bottom half of the word lines. More generally, the keys and inverse pairs could be interleaved in any desired fashion, as long as it was consistent for all of the keys in the block; however, this would require keeping track of the arrangement. The interleaved pattern along the NAND chain may be preferred since the data can be inversely program in another WL without loading the data again. There are some other coupling effect may also benefit from interleaving the inverted and non-inverted data on adjacent word lines. In terms of performance for this type of embodiment, for a 16KB page of 64 bit keys, if a duplicate key/inverted key pair is kept, this is 8KB, or 64,000 keys. At 35us per sensing, this gives 1.82C/s/plane. If 4 planes are operated in parallel, this is 7.3 CG/s at around 200mW.

[0061] For either of the embodiments of Figure 4 or Figure 6A, the method uses the inherent “AND” functionality available in a NAND Flash memory to compare thousands of keys in a single sensing operation. This method has several major advantages over traditional CPU- or semiconductor-based CAM memories. For one, as the comparison is done “on die”, there is no need to transfer the data out of the memory. This saves both time and IO power. Furthermore the actual comparison operations use

less power than conventional memories. As all of the bit lines are sensed at the same time, with only the matching NAND chain is conducting current, the NAND based CAM is highly parallel; for example, in a NAND flash memory with 4 x 8KB planes, $(32K \times 8\text{bits/byte})/2 = 128K$ keys can be checked in one sense per die. If a sense can be done in 35us, an even/odd sense as described above with respect to Figure 4 will take 50us. This is 128K keys in 50us, so that an entire 8GB die (2000 blocks) could be sensed in ~100ms. The corresponding energy consumption is on the order of 200mW. To increase performance, multiple die can be operated in parallel.

[0062] As noted in the Background section, keys can be stored in a CAM as either sorted, in which case a binary search can be used; or as unsorted, in which case a linear search is used. This is also true of a NAND based CAM, except that as NAND based CAM can be searched at the block level, in a sorted CAM the keys need only be sorted to the granularity of the block or the number of blocks that are sensed in parallel. The CAM allows for a binary search, but at the block level due to this parallelism. Even for linear searches, this degree of parallelism can make linear searching comparable or even faster than binary searches for fairly large data sets. Again, for any of these arrangements, performance here can also be improved by running multiple die in parallel.

[0063] The keys can be sorted based on a given number of most (or least) significant bits. A sorting based on significant bits is generally most useful when the key or content being searched is not a hash value, but a set of characteristics or data itself. In this case, the sorted data in each block would all share a certain number of most significant bits for their keys.

[0064] Content addressable memory exist in both binary form, where the search key consists of 0s and 1s as described above, and ternary form, where the search key can also include “don’t care” value. As discussed above, when a high read value is broadcast along a word line, all of the cells along that word line will conduct regardless of its state. This property allows for a “don’t care” value to be implemented by setting the corresponding word line to the high read voltage for both the key and its inverse; that is, when sensing with the key and its inverse (in either the second read of Figure 4, or the lower half of the word lines), the don’t care values are set to the high read value for both the key and its inverse, while the other values of the key are inverted as before.

[0065] These properties of a NAND based CAM also make it particularly suited to a number of other uses. For instance, as large numbers of keys can be searched in parallel, this allows for all copies of the same key in the searched blocks to be determined in the process, improving efficiency of de-duplication operations of the sort that are valuable in cleaning up data bases. The NAND structure also makes for a CAM useful as a Bloom filter as an intersection of multiple search keys can be formed by setting any values that differ between the keys to the high read voltage in the combined search key, which can then be used to search the horizontally stored keys of one or more blocks in parallel.

[0066] The use of “don’t care” values can also be used to perform a type of “iterative” search. This can be used the keys may have, or possibly have, some number of bit errors. One or a series of reduced keys, where some subset of values are set to “don’t care” for different entries, are used for searching. Several different such keys, where a different subset of values is masked out in each, are then iteratively used to check for matches among the written keys.

[0067] Another example of where a series of reduced search keys can be employed is where the content is itself a data set, as opposed to say a hash value. For example, it could be values from image processing. In such a case, the content of the block could be searched to a desired number of significant bits, by setting bits of lower significance to “don’t care”. Similar arrangement could also be used for patterning matching of the content or for cases where the keys are properties of main data content.

[0068] Consequently, due to its parallelism, relatively low power consumption, or both, NAND based CAM can be used in many applications, such as data base searching, voice recognition, DNA matching/genome searches, cryptography and so on. It can lend itself to CAM based indexing and can be incorporated, for example into CAM indexed SSD systems.

[0069] So far the discussion has mainly considered the case of binary NAND memory for CAM use. More generally, multi-state (MLC) memory can also be used; for example, in a mixed binary-MLC memory, the keys could be stored in binary memory for CAM use, while data to which the keys pointed could be stored in MLC areas. It is also possible to use MLC NAND memory for CAM, using 2 to 3 bits per cell, for

example, in key matching. Using 2 to 3 bits per cell, the NAND chain can store longer keys. In the sort of embodiment described with respect to Figure 6A, where a key and its inverse are written into the same word line, a 128 cell NAND chain in binary operation can store 64 bit keys, while a 128 NAND chain with 2-bits per cell can store 128 bits keys. Similarly, 3-bits per cell operation can store 192 bit keys.

[0070] Figure 7 shows an exemplary encoding of 2-bits per cells for four state memory cell operation. As shown, the erased state is encoded as 11, the first state up (or “a” state) is 10, followed by 00 (for the “b” state) and 01 (or “c” state). The various sensing levels are also shown. Figure 8 shows how the data states and the complementary data used for the inverted keys correspond. Figure 9 shows an example of how a key ({00111001}) would be encoded onto a 4 cell NAND string on bit line BL and its complement on bit line BLB. In the MLC CAM embodiment, the system can use one or two word lines along the NAND chains to store the parity bits of each NAND chain in order to check on the integrity of the NAND chain. As with the binary case, manufacture defective columns can be isolated out and more redundancy along the word lines (duplicated data) can further protect the keys’ integrity. Also, the complementary data shifted as illustrated in the figures to provide more sensing margins.

[0071] In the above, keys were written down the bit lines of the array, with the search key broadcast along the word lines, allowing the keys along a block’s bit lines to be searched in parallel. The arrangement can also be reversed, where NAND array can also be operated so that the content or key matching is in the word line direction. Under this arrangement, one or more keys would be written along each word line (that can be very long keys), an arrangement that can be useful in several different circumstances. Multiple short keys can be stored along the word line direction as well. If the keys are encoded so as to have significance as a 2D array of values, this would allow for content searching in both of the bit line and word line directions, although the more typical situation would just be for content matching in the word line direction. For instance, as word lines are much long than bit lines, a word line based CAM allows for the use of longer keys. Also, as data is written in page along word lines, it may be more convenient, at least initially, to write incoming key data along word lines. This would then allow for key to be searched as written along the word lines. If desired, the keys

could then be rewritten along bit lines, where they could then be searched as described above. Due to the amount of data involved, to move the keys from a word line orientation onto a bit line orientation, once a block worth of the keys are assembled they could be read off of the NAND memory into the controller and then rearranged and written along bit lines; alternately, this rotation of the keys could be performed on the NAND device, although this will typically require the memory chip's structure to be modified to accomplish this. This process of initially writing the keys along word lines, transferring the keys to the controller, and rearranging them to be rewritten onto bit lines can also include the sort of sorting process described above prior to rewriting them back in a bit line orientation.

[0072] The process of matching of content in word line direction is illustrated with respect to Figure 10. As the system receives keys, these can be formed into pages of one or more keys and written into the memory array 901 along word lines. To match content, the system inputs the matching content of one or more search keys into a matching buffer or register 905, which can then be used to look for duplication content along the word line. The data along a word line is read from memory array 901 into a buffer or register 903. The memory can then perform internal match operations between the read data in buffer 903 and search data in buffer 905, where some number of bits ignored, if desired. The ignored bits can either be to "don't care" values, because some read error can occur on the read. These ignored bits can be corrected with error correction in the controller when a match is found. Typically, the smallest length of key/content along the word line that can be compared is 1KB, while the longest length of key/content that can be compared in one plane is 16KB. If the key length is smaller than 1KB, the key can be duplicated in chunks patterns to do the pattern matching with more parallelism. Then the matched case will produce a group of "1" and the un-matched case will produce 50% "1"s. Circuitry can detect if a word is all "1"s to judge the match or miss. If there are some "0"s in a word, this word can be discarded as a miss. In order to prevent that some read errors produce a miss when it should be a match, a majority voting circuitry can be employed to choose the word with majority "1"s for matching. Some words can be masked out by marking the isolation latch to be "ignored". To simplify operations, it is typically preferable to write the beginning of a file to aligned with certain columns. After finishing a compare on one word line, the next word line content can be compared in a similar sequence.

Key-Value Addressed Storage Drive Using NAND Flash Based CAM

[0073] This section considers the incorporation of the sort of NAND Flash content addressable memory (CAM) or content addressable storage (CAS) described above in a storage drive. Conventional storage drives, such as solid state drives or hard-disk drives (SSD or HDD), are addressed by a logical block address (LBA) to read and write data stored on their physical media. These employ logical to physical address translation tables to locate the data, where the address translation table is stored on flash, in DRAM, or on magnetic media and is updated on the basis of sectors, bytes, or pages. Typical sizes for such addresses are 32, 48, or 64-bits. In some applications, such as in databases, it is desirable to have large keys (of hundreds or thousands of bits) that can address a space much larger than the number of elements in the database. In these cases a content addressable memory utilizing key-value pairs is used to index the elements stored in the device.

[0074] In a content addressable memory, data is written as a key-data pair. To retrieve the data, a search key is supplied; all the keys in the memory are searched for a match. If a match is found, the corresponding data is retrieved. This section presents a storage drive using a Flash based NAND array as described in the preceding section as a content addressable memory that is addressed using key-value pairs instead of a logical block address. This drive can provide both Binary and Ternary search capability, meaning that bit patterns in the key can have the values 1 or 0 as well as “don't care” entries. This type of NAND based CAS drive can then be used to replace other implementations of CAM or CAS functionality, such as those employing a database, that would usually include a host CPU, DRAM, and storage media.

[0075] Consequently, this section applies the of operation of a NAND flash memory as a pattern matching engine from the last section to a storage device that is indexed using key-value pairs instead of conventional logical block addresses. The device can use a standard transport protocol such as PCI-E, SAS, SATA, eMMC, SCSI, and so on. When used in a pattern matching mode, the NAND cells not only store values, but can also be used to compare their stored values with an input value. In the examples of this section, target patterns are stored along bit lines, although the word line based storage discussed above can also be used. In the bit line example, the pattern to be matched is broadcast down word lines. If all the elements in the NAND chain match their target

pattern, the NAND chain (bit line) will conduct. The position of the conducting bit line can be used as an index in to another table that can be used to retrieve data that is associated with the target key. This is shown in Figure 11, that expands upon Figure 1.

[0076] In Figure 11, the bit lines BL0, BL1, ... run down the columns of the array and are written with corresponding keys Key 0, Key 1, ... as previously described. The word lines are then biased according to the search key (here Key 2) so that it is broad to all of the bit lines spanned by the word lines. When a matching key or keys are found, the column address of the bit line is then input as an index to find the data set, also stored on the drive. A number of arrangements are possible, where, for example, the keys could be stored in binary or MLC arrays optimized for CAM use, while the data is stored in more standard MLC arrays.

[0077] A drive using such a mechanism can then be used to search for key-value pairs in a large search space, perform general pattern matching (using bloom filters), or be used for determining set membership. Some of the advantages of a drive using such a scheme include low power usage and high bandwidth. As data does not need to be moved from the NAND array to a separate computational module for comparison, power consumed on IO operations is reduced. Furthermore, since only bit lines that match a given search pattern will conduct, the NAND comparison operation is also low power. With respect to bandwidth, a single NAND die is capable of doing, say, 256K 64-bit comparisons in under 50 us, working out to under 200ps per comparison. Additionally, multiple die can be operated in parallel to increase bandwidth or to increase the effective key-length. Potentially 8Gb (~8G keys) of 64-bit keys can be searched in ~100ms in a single die based on current design.

[0078] The idea of using a storage drive with a common interface, such as SAS, PCI-E, SATA, and so on, but which is written using key-value pairs instead of a conventional logical addressing is consider further with respect to Figure 12. In the write method, the host will write Key-Value pair (K, V) to the drive. The drive will store the Value V in a data store at some address in the Data table of Figure 12, as illustrate at (1). The drive will store the key value K on a bit line “i” in a block of an array of the drive, as shown at (2) of Figure 12. The drive will make an entry in the block table at address i, with a pointer to the value V, as shown at (3). The column address which has the matching key can be output from NAND memory from status bits.

[0079] In NAND flash memory, data is written in word line based pages. Because of this, as previously discussed, the keys may be initially written along word lines, then rearranged to be written along bit lines, or first stored in RAM and then sorted into bit lined oriented keys. (It could also be possible for the host to have already taken care of this bit line based orientation for the keys, although it will generally be preferable for this operation to be transparent as seen from outside of the drive, with a host just providing basic key-value pairs and not having to engage in such data manipulations.) The controller will take care of assigning the keys and values to physical addresses and of determining the needed addressing structures to translate the key into the corresponding data location. The key to value mapping tables can be maintained in much the same way as the usual logical to physical mapping tables as far storing them and updating them, such as mappings using look up tables or based a correspondence formula. For example, the column address can be mapped to metadata in the primary storage flash management layers.

[0080] In a variation, if the drive itself has a key generating ability, such as a hashing algorithm using by the controller, just the data set itself could be sent to the drive and the corresponding keys generated on the drive. This would also allow a different host without the key generating ability to send data to the drive, where the key-value then be generated and stored. Under this sort of arrangement, for a host to access a value using a key, the host would need to use the same key generating algorithm (such as from a Secure Hash Algorithm (SHA), for example) as being used by the drive.

[0081] Looking now at the read process, the host will send the drive a key (K) that is then used to search key blocks. As discussed in the previous section, the key blocks may be sorted, in which case a binary search can be used; or they can be unsorted, in which case a linear search is used. For each key block, the drive will apply the key K to the word lines. If a matching key exists along a bit line in the block, NAND flash will register a “1” at the bit position “j” associated with the matching key. The value “j” can then be used as an index to the associated block table, as represented at (4) in Figure 12, to retrieve a pointer, (3), to the associated value V in the Data Table. If all key blocks are searched without finding a match, the drive can return an “element not found status” or error.

[0082] The CAM NAND can be incorporate into the same memory system as that in

which the associated data is stored, such as an SSD, in which case the data corresponding to the search key can be provided directly to the host. In other embodiments, the CAM NAND could be a separate device used to provide the sort of CAM-based operations described here, while the associated data could be stored separately, in which case as address or other pointer to the corresponding data on the separated device would be provided.

[0083] Consequently, storage drive can benefit by use of the inherent “AND” functionality available in a NAND Flash memory to compare thousands of keys in a single sensing operation. A storage drive of this type has several major advantages over traditional CPU- or semiconductor-based CAM memories. First, because the key comparison is done “on die”, there is no need to transfer the data out of the memory. This saves both time and IO Power. Furthermore the actual comparison operations use less power than conventional memories.

[0084] In addition to the low power and high bandwidth advantages discussed in more detail earlier, this scheme has the advantage that write times can be shorter if data is searched in a linear mode. Most databases spend time and energy sorting and maintaining tables to enable fast, binary type, search capability for when data is read. In one embodiment of this invention, the writes of data and keys can be done in a random fashion making writes times of $O(1)$ complexity. Searches will use a linear search mechanism which is highly parallelized but is still $O(N)$ complexity. This is less efficient than the $O(\log N)$ of most binary searches and is a tradeoff between insertion time vs. lookup time. Also, it should be noted that even when data is stored in a sorted manner, the high degree of parallelism in searching mean that the sorting only needs to be done to the level at which the search is done, namely to the granularity of block or number of blocks searchable in parallel.

[0085] The sort of NAND flash base CAS drives can be applied to a number of applications, including those described in the previous section. One set of examples of these exemplary applications is for de-duplication using pattern matching (CAM) NAND to store the hash keys. Incoming data can be sent through the hash function to generate the content related fingerprints. The fingerprints can then be searched with the existing hash keys to see whether the data already exists in the data storage. If it does already exist, no write action is taken; but if the data does not yet exist, then the new data

will be written into the storage. The de-duplication can be done when the data is backing up, during garbage collection operations of the primary storage, or in-line as the data comes in from host. Another application is for virtual memory management, which can be done similarly to de-duplication. The drive can also be applied to the Human Genome, where the drives stores signatures in the CAM NAND so that any segment of the DNA sequence can be searched. The drive also lends itself to parallel computing where, a mathematical NAND function can be done inside of the NAND memory.

[0086] As noted above, the CAM NAND operation has the keys oriented along bit line, whereas NAND memory written along word lines. Consequently, as the keys come in from a host, they need to be accumulated in a buffer memory of some sort, transposed to a bit line orientation, formed into pages (including adding any inverse keys as needed), and transferred to the NAND device for writing. This is illustrated schematically in Figure 13.

[0087] As shown in Figure 13, a host 1301 (a PC, for example) can take the data files and generate the corresponding keys, such as using a Secure Hash Algorithm (SHA) to generate a 64 bit hash key, which can then be transferred over to a buffer memory 1303 on the memory system, where the keys can be accumulated. The transposing buffer memory 1303 is used to align the keys for writing in the NAND CAM memory 1305. Once a sufficient number of keys, say 4MB keys for a NAND memory 1305 with 4MB blocks, the data can be transferred over for programming as pages along the word lines. Figures 14 and 15 give some examples of hardware implementations for the transposing buffer memory, but this can be implemented in various other ways, such as by use of a field programmable gate array (FPGA). For example, a blocks worth of keys could be accumulated in an FPGA and then read out a word line at a time and transferred over to the CAM NAND for writing.

[0088] Figure 14 is a schematic illustration of a hardware implementation of the transposing memory in FIFO style. The data can come in as, say, 64 bits keys or indices and is saved in column-oriented 64 bits registers. The registers are chained into a FIFO arrangement so that when a new key comes in, the previous keys shift over by one column to the right. Once there are enough keys for one page (i.e., 8KB to 16KB for a typical design) to be programmed into the CAM NAND memory (or early, if smaller

number of keys are to be used), the pages are shifted over to the NAND for programming into the array there. Under this arrangement, if desired, the keys may be searched while still in the FIFO before being programmed, as the keys can be shifted out one at a time for comparison.

[0089] Figure 15 is a schematic illustration of another hardware implementation for transposing the data keys using more of a RAM style arrangement. The data can come in as, for example, 64 bit keys or indices and be saved in 64 bits registers, being accumulated in a relatively small, 16x64 array 1509 in latches. The 16x64 bits of the small array 1509 can then be shifted over a bus 1507 a word (16 bits) at a time into the RAM 1503. After shifting the 64 words into RAM, the small array 1509 can accumulate next 16x64 bits. This process can continue until the RAM 1503 is full or it is otherwise desired to write in the keys, at which point is programmed in the CAM NAND memory. Under the arrangement of Figure 15, if it is desired to search the keys before they are written into the CAM NAND, another RAM buffer storing the keys without transpose can be kept for this search purpose.

Applications to Data Analytics

[0090] The sort of highly parallel operations using a memory device of a NAND structure as a content addressable memory described in the preceding sections can also be applied to performing data analytics. This allows for massively parallel computing to be applied to various analytic applications, where the computing be performed inside of the storage and remotely from the server. This arrangement can also allow processing to be done in real time, using inline processing, and also allow for the analytics to be executed without input/output transmission limitations. Consequently, these techniques and structures can be applied to many applications, from crunching large amounts of data in data warehousing applications, quantitative analysis of financial data, and other data analysis intensive uses.

[0091] A number of topologies can be used, with one example shown in Figure 16. In this example, a memory system 1601 is a computing solid state drive (SSD) that includes a main storage SSD section 1603, the NAND device used can be normal NAND devices as well as CAM type NAND. In the exemplary embodiment, the NAND portion 1605 as again taken as an EEPROM based flash memory when a specific concrete example is needed. A host 1611, such as a PC or even a network

connection, provides data and any instructions for analytics to perform on the data to the memory system 1601. The data can be supplied to the NAND section 1605 to be stored for analyzing and then to the main storage section 1603, allowing for in-line analysis if desired, or stored directly in the main storage section 1603 and retrieved to NAND module 1605 when analysis is requested. In some cases, such as when the data are key/associated data pairs, the keys could be stored on the NAND 1605 and the associated data going to the main storage section 1603, where the system can maintain a key-data correspondence as described in preceding sections. In case where, in response to a request from the host, data is transferred from the main storage section 1603 to the NAND section 1605 to have analytics performed, the CPU or GPU or SSD controller could be used to perform some initial manipulations (choosing data subsets, generating hash values, and so on) as needed before writing the data into the NAND structure of 1605.

[0092] The arrangement of Figure 16 will be used as the exemplary embodiment in the following, but a number of variations or other topologies can be used. For example, the main data storage section 1603 need not be a SSD, but could be hard drives or other data storage. Also, the NAND portion 1605 need not be incorporated into the same system as the main storage 1603, but a separate system for this portion used in conjunction with a bulk data storage system. In other cases, where the amounts of data are manageable by the CAM-based NAND system itself, the NAND system can be used directly with the host for performing data analytics. For instance, a portable device incorporating the CAM NAND and some additional flash storage may be sufficient.

[0093] Figures 17-20 illustrate how the NAND array can be used to perform analytics in parallel for all of the columns of the array when the data includes both categorical (i.e., data that can fit into multiple categories, such as (red, blue, green) or (yes, no)) data as well as numerical range detection. Due to the CAM nature of the memory described here, multiple categories can be handled. In this example, categorical and numerical data can be stored along the same NAND strings, but the categorical data is saved in a binary format, while the numerical data can be save as binary (D1), 2-bit per cell (D2), 3-bit per cell (D3) or other multi-state format. It should also be noted that the categorical/numerical distinction is not necessarily hard and fast, as the techniques

described here allow for the processing of numerical data to make it into categorical data for purposes of analysis as, in some case this can be faster than performing numerical comparisons. The more bits per cell, the fewer the number of word lines that will be used to store the data, but with the increased complexity involved in such multi-state operations. The analytics will generate a match for the specific query and the match results can be counted inside the NAND or outside of NAND. As discussed further below with respect to Figures 30 and 31, the counting can be done inside NAND digitally, which is precise, or in an analog, which is faster but less accurate. When counting outside NAND, the match results will be transferred to controller and the number of “1” or “0” will be counted there.

[0094] Figure 17, at left, shows how two of the data sets of a block are written onto NAND strings along bit lines BL_n and BL_m. At upper portion of each string is entered categorical data in binary form, with some numerical data further down, where a 3-bit per cell format is shown, the bits arranged top to bottom as least to most significant bits. Here the shown categorical data word lines can be searched first, with the analytics then performed on numerical data for the matched “category”. The categorical data can be “don’t care” or not written with the numerical data at the same memory block. The numerical data can then be sequentially analyzed, here starting with the most significant bit, by reading one word line at a time, placing the appropriate read level (CGRV) on the MSB word line. At right is a schematic representation of how the numerical data can be compared as a sequence of queries. Each of the bit lines has an associated set of latches that can be used to keep track of the results of the sequence of analytic operation, where an example of how the results are assigned to the latches is shown in Figure 18. (The data latches here are labeled XDL, UDL, and, further down, LDL for transfer, upper, and lower data latch respectively, to correspond to the arrangement such as that described in US patents numbers 7,206,230 and 8,102,705, where more detail on such latch structures can be found, and also see Figures 28 and 29 below.)

[0095] The table of Figure 19 shows an example of a compare to see whether a number is greater than 010011001 for four data values. The search is here done from the most significant bit down towards the least. When the most significant bit (MSB₉) is searched, the top value is found to be greater than the search number, the latches are set

accordingly, subsequent reads are ignored and no updates are made. For the other three values, the results are indeterminate. When the next bit, MSB8, is checked, the second data is still indeterminate, but the lower two values are found to be less than the search values so that the latches are set accordingly and no more updates are needed afterwards. Continuing with the second data value, the MSB7 result would again be indeterminate and is not shown, but the MSB value establishes that it is greater than the search values and the latches are set accordingly. At this point, the final search values for this data set are established, as shown in the right hand most column. If there were still indeterminate data, the process would continue on towards the least significant bit until the final search results were all established. The final results will be shown on one of the data latches, such as XDL=1 in this case. The match to fit the query can be counted later or saved to another word line for further analysis in combination with other query criteria.

[0096] Figure 20 is an example of another search to see which data values are between 123 and 231. For the first read, the first digit of the data values are checked against the upper bound, which is found to have exceeded for the first number, putting it out of the range so that any subsequent reads can be ignored. The second number is found to equal the MSB upper bound, with the bottom data to be under the MSB upper bound. At the second read, the second digit of the second number is found to exceed the upper bound, so the latches are set and no further updates are needed. For the fourth data, the second read finds this below the lower MSB values and, consequently, outside the range so that the latches are again set and no further updates needed. The second read for the third row data finds it to equal the MSB of the lower bound, so that the next digit is checked against the second digit of the search's upper value in the third read and the second digit of the search's lower value in the fourth read. The final search result is then shown in the far right column.

[0097] Figures 21 and 22 illustrate how to perform to maximum and minimum value searches. In Figure 21 for the max search, the process illustrated for 5 NAND strings each with 9 bits arranged with LSB to MSB top to bottom. At left, the MSB row is searched and loaded into the UDL latch. In this example, the left two most NAND strings have a "1" for the most significant bit. The other columns can be ignored for the rest of the search. The process works its way up the rows, where the next two most

MSBs are indeterminate. In the fourth row up, the two left columns are different, where the results as loaded into LDL show that the second column has the max value. (Here, two latches are shown for illustrative purposes, but this process can be done in a single latch per bit line by overwriting as the row is read one at a time.) Figure 22 similarly illustrates a min search, again working its way up from the MSB. At left is shown the situation after working up to the fifth most significant bit, where the outermost column have both had zeros up until that point, as reflected in UDL. At right of Figure 22 shows the result of two reads later as loaded into LDL, showing the left most column to be the minimum.

[0098] Max and min search can be performed on file size. For a max, the memory can find the file size with most number of digits along the NAND chain, then find the next largest files by eliminating the small numbers. For a min, the memory can find the file size with least number of digits along the NAND chain, and then search for the next smallest files by eliminating the larger numbers. Parts of a file system can be stored in this manner.

[0100] The array structure allows for the data analytics to be done one row at a time, as they can be done by reading one word line at a time. For example, the array structure can also be used to perform arithmetical operation, such as addition, subtraction and multiplication, on the numeral data along the NAND strings. The process is schematically illustrated in Figure 23 for a summing operation. Here the data sets of block N 2301 can added to the corresponding data sets of block M 2301 on a row by row basis for each bit line. The result can then be written into a SUM block 2305. Here, block N has NAND strings N1 to N128K, block M has NAND strings M1 to M128K, and the SUM block similarly has NAND strings SUM1 to SUM128K. Selected numerical data values can then be added by reading the values of each word line from NAND string i of blocks M 2301 and N 2303 into the associated latches 2307 a word line at a time, added there, are written back to the SUM block from the latches, $SUM_i = N_i + M_i$. If the NAND has more data latches, more bits can be accumulated before writing to the SUM block. In the case of 3 data latches, there are 4 numbers can be added before writing to the SUM block. In the case of 5 data latches, 16 numbers can be added before writing to SUM block.

[0101] In the example of Figure 23, and also in those discussed below, the NAND array is organized as to have a block structure, such as found in flash memory. More generally, the word lines can be any of the word lines of the array; and when the memory has a block structure, these need not be from different blocks, as, for example, when adding two different word lines of numerical data from the same data set. Also, more than two data can be processed and saved into the data latches before writing the result to a new word line. For example, with 3 data latches the memory can add 4 pages before writing to a word line, saving the carry in the data latches for the next bit addition. For 5 data latches, it can add 16 pages and then write to different word line once, and so on.

[0102] Figure 24 illustrates an example of how the latches can be used in this process for addition. The latches associated with each bit line are here labeled as UDL, LDL, and XDL. Figure 24 illustrates a single one of each of these latches with values read from different word lines holding different values as these are read sequentially for a 13 bit number from LSB to MSB. (That is, the arrow represents the sequence of reads in time or, equivalently, down word lines for a single UDL, LDL, XDL set of latches associated with a single word line. UDL contains the data set collected at time A from a first block (such as block N of Figure 23) and LDL contains the data set collected at time B from a second block (such block M). The XDL latch can hold any carry bits. The two data sets can be added and stored back in LDL, with the summation then programmed back into another block. Other operations (multiplication, subtraction, division, etc.) can be similarly performed: for example, subtraction can be done as one data added to the complement of the other data. In addition, floating point operations can similarly be performed by properly aligning the digital information so that the points align for the operands.

[0103] Figures 25A-C give some more detail on the mechanics of some arithmetic operations as these are execute in the latches. Figure 25A looks at addition, specifically $10+3$, as noted at top. Here “10”, or 1010 in binary, is read from block A and “3” (0011 binary) is read from a block B. These number are shown listed for blocks A and B in the left column, MSB to LSB written bottom to top. At times T0, T1, T2, and T3 these are read into a the latches UDL and LDL, with the carry being held in the XDL latch, as

shown underneath. The results are written back into Block C from the values shown latched there.

[0104] Figure 25B illustrates how to perform subtraction of two numbers $N1$ and $N2$ to form the difference $N1-N2$. This is done by adding $N1$ to the 2's complement of $N2$ plus 1. A specific example, here again using 10 and 3 to determine $10-3$ in the latch structure is shown: in the top row is the binary form of 10, in the second row the 2's complement of 3 plus 1 ($3c+1$), and the result is shown at bottom. Any overflow bits need to be discarded, with the result being the binary form of 7. Figure 25C shows how multiplication can be done using bit shift and addition, where 10 and 3 are again used as the inputs.

[0105] Figures 26A and 26B look at examples of where, in addition to the XDL latches there are additional latches available on each bit line beyond UDL and LDL, such as is found in a multi-state memory device. Figure 26A looks at the case of 3 data latches. As shown, data from four blocks (or, more generally, four pages) are written in. This allows for four numbers to be added or subtracted in a single write. Figure 26B shows a 5 latch case, allowing for up to 16 numbers to be added or subtracted in one write.

[0106] Some examples of an application to financial data analysis are illustrated with respect to Figures 27 and 28. Figure 27 illustrates an example of loading the stock data for a single stock historical data analysis, where for the 128k (in this example) bit lines can each be assigned to given stock or other financial instrument, with day for each day written to a different block. The data for each stock can be lined up along the NAND strings. With each stock taking a bit line, for an array of 128K bit lines a total 128,000 stocks can be evaluated simultaneously. The price per day can then take different blocks or different locations of the NAND chain. Using the blocks for different days, operation such as averages, linear regression, and so on can be performed using the data from the corresponding blocks, where the analyzed data can be saved in a new block.

[0107] To perform other types of analyses, the data can be arranged differently on the array. An example is shown in Figure 28 where the data sets are arranged to perform a correlation study. In a first block the data for up to 128K stocks on a given day are entered. In the second block, the data from different stock B pre-processed from same chip or different chip will align up with the pre-processed data for stock A on the same

bit line. Then the correlation between stock A and B can be calculated accordingly. Depending on the complexity of the desired operations and the on-chip ability given the NAND device, these operations can be performed on chip or with the help of the controller. For example, if needed for particular applications, the NAND can include fairly complex, but specific, operations.

[0108] Figures 29-31 consider data arrangement for analytics in more detail. So far, the discussion of this and the preceding sections have largely considered the data sets or keys being analyzed on the NAND array as being on a single NAND string. However, more generally, as each bit line can have many NAND strings formed along it that are connectable to the same data latches, this allows for data (a schema) to be arranged in few separate blocks along the same NAND bit line. This is illustrated schematically in Figure 29, where some numeric data of a data set, arranged from most to least significant bit, is stored on the same bit line, but in NAND strings from separate blocks. (Here these are shown as adjacent, but need not be so in the more general case.) The bit line then is connectable to the transfer data latch XDL and, through XS W (transistor switch), the corresponding sense amp (SA) and data latches UDL, LDL, where these can correspond to upper and lower page data latches in a more typical multi-state operation. (More detail on the latch structure described in Figure 30 and, in the next paragraph, Figure 31 can be found in US patents numbers 7,206,230 and 8,102,705.)

[0109] Data can also be arranged inside a group of (typically adjacent) bits lines that share the same data bus to communicate with a set of data latches. For example, a set of 8 or 16 adjacent bit lines with such shared structure could store each data set on multiple ones of these bit line groups. Figure 31 schematically illustrates a latch structure where 8 NAND bit lines can process the data with shared SBUS and DBUS through logic operations from YBOX circuitry, so that a schema can be stored in up to 8 associated bit lines sharing same data latch set.

[0110] The results of the analytic operations can be computed according to various data counting methods. As illustrated schematically in Figure 32, the counting can be done digitally inside the CPU, toggling the data out to RAM for counting. Digital counting can also be performed inside the NAND device, such as by binary search or shooting chain. Analog counting can also be done inside the NAND, which, while less accurate can be done more quickly. Figure 33 shows some elements of such circuitry for

counting quickly with analog wired OR circuitry: here, the data is applied to the gates of a set of transistor connected in parallel, each connected in series with a transistor controlled by an analog bias. The transistors are fed by one leg of a current mirror, the other leg of which is connected to ground through a transistor acting as acting to set a digitized current level.

[0111] Figure 34 illustrates how analytics can be applied to large file systems. As file systems grow, there is an increasing demand to perform complex queries, such as “How many files have been updated since 10 days?” and “Which are the top five largest files that belong to John?” The first is an example of aggregate queries which provide a high-level summary of all or part of the file system, while the second is top-k queries which locate the k files and/or directories that have the highest score according to a scoring function. The incorporation of the NAND section as described with respect to Figure 16 provides a simple solution for performing such queries.

[0112] At left, Figure 34 shows a file system. The basic file data (owner, time stamp, file size, etc.) is saved into NAND in vertical NAND chains as shown at the right of Figure 34. Performing the analytics by the NAND SSD in this way saves the trouble needed to build the data tree structure as shown in the file structure on the left. Then aggregate queries can, for example, be searched on the time stamp that can be located in a few word lines against a certain date. Top-k queries can be done, for example, by identifying “John” and the file size. (As noted above, min and max searches can also be done for file sizes, allowing for portions of file systems to be stored on this basis.)

Conclusion

[0113] The foregoing detailed description has been presented for purposes of illustration and description. It is not intended to be exhaustive or limiting to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. The described embodiments were chosen in order to best explain the principles of the disclosed technology and its practical application, to thereby enable others skilled in the art to best utilize the technology in various embodiments and with various modifications as are suited to the particular use contemplated.

IT IS CLAIMED

1. A method of analyzing data, comprising:
writing a plurality of data sets selected by a host in a memory array of a NAND architecture of NAND strings formed along bit lines and spanned by word lines, where the data sets are written oriented along the bit lines with corresponding data elements arranged to lie along the same word lines;
concurrently biasing a plurality of word lines to a first set of sensing voltage levels corresponding to host determined parameter values; and
determining those of the bit lines that conduct in response to said biasing of a plurality of word lines to the first set of sensing voltage levels.
2. The method of claim 1, where the selected data sets are provided from the host.
3. The method of claim 1, where the selected data sets are selected by the host from a memory system.
4. The method of claim 3, wherein the memory system is a solid state drive (SSD).
5. The method of claim 3, wherein the memory array is part of the memory system.
6. The method of claim 1, where in the sets of data include data of a first variety and a second variety, the method further comprising:
subsequently concurrently biasing a plurality of word lines to a second set of sensing voltage levels corresponding to host determined parameter values; and
determining those of the bit lines that conduct in response to said biasing of a plurality of word lines to the second set of sensing voltage levels,
wherein the first and second sets of sensing voltages are the same for data of the first variety, but differ for data of the second variety.
7. The method of claim 6, wherein the first variety of data is categorical data and the second variety of data is numerical data.

8. The method of claim 7, wherein the first variety of data is written in a binary format and the second variety of data is written in a multi-state format.

9. The method of claim 6, wherein the first and second sets of sensing voltages each include on word line biased to a data dependent value for data of the second variety, with the word lines of the other word lines for data of the second variety being biased to a data independent read value.

10. The method of claim 1, further comprising:
subsequently determining a second set of sensing voltage levels dependent upon the result of said determining;
concurrently biasing a plurality of word lines to the second set of sensing voltage levels; and
determining those of the bit lines that conduct in response to said biasing of a plurality of word lines to the second set of sensing voltage levels.

11. The method of claim 10, wherein the memory array is formed on a first memory circuit including sensing circuitry connected to the bit lines, the sensing circuitry including a set of data latches corresponding to each of the bit lines, the method further comprising:
storing the results of determinings in response to the biasing to the first and second sensing voltage levels in the data latches; and
performing with the data latches one or more logical/arithmetical operations upon the results of the determinings in response to the biasing to the first and second sensing voltage levels.

12. The method of claim 11, wherein the logical/arithmetical operations include arithmetical operations performed digitally.

13. The method of claim 12, wherein the logical/arithmetical operations include arithmetical operations performed digitally on the first memory circuit.

14. The method of claim 12, wherein the logical/arithmetical operations arithmetical operations performed digitally on a controller circuit that is part of a memory system including the first memory circuit.

15. The method of claim 11, wherein the logical/arithmetical operations include arithmetical operations performed using analog techniques.

16. The method of claim 15, wherein the logical/arithmetical operations arithmetical operations performed using analog techniques on the first memory circuit.

17. The method of claim 16, wherein the analog techniques used to perform said logical/arithmetical operations include determining the amount of current drawn by a number of transistors connected in parallel, the number of said transistors depending upon the result of the logical/arithmetical operations.

18. The method of claim 11, further comprising:
providing the results of the logical/arithmetic operations to the host.

19. The method of claim 11, further comprising:
storing the results of the logical/arithmetic operations in non-volatile memory;
and
subsequently performing one or more analytical operation upon the data.

20. The method of claim 19, wherein the subsequent analytical operations are performed on data including one or more of the data sets, the method further comprising:

performing one or more logical/arithmetical operations upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations; and

providing to the host the results of the logical/arithmetic operations based upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations.

21. The method of claim 19, wherein the subsequent analytical operations are performed on data including one or more additional data sets, the method further comprising:

performing one or more logical/arithmetical operations upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations; and

providing to the host the results of the logical/arithmetic operations based upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations.

22. The method of claim 1, wherein the memory array is formed on a first memory circuit including sensing circuitry connected to the bit lines, the sensing circuitry including a set of data latches corresponding to each of the bit lines, the method further comprising:

storing the result of the determining the data latches; and

performing with the data latches one or more logical/arithmetical operations upon the results of the determining.

23. The method of claim 22, further comprising:

providing the results of the logical/arithmetic operations to the host.

24. The method of claim 22, further comprising:

storing the results of the logical/arithmetic operations in non-volatile memory;

and

subsequently performing one or more analytical operation upon the data.

25. The method of claim 24, wherein the subsequent analytical operations are performed on data including one or more of the data sets, the method further comprising:

performing one or more logical/arithmetical operations upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations; and

providing to the host the results of the logical/arithmetic operations based upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations.

26. The method of claim 24, wherein the subsequent analytical operations are performed on data including one or more additional data sets, the method further comprising:

performing one or more logical/arithmetical operations upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations; and

providing to the host the results of the logical/arithmetic operations based upon the result of the subsequent analytical operations and the stored result of the logical/arithmetic operations.

27. The method of claim 1, wherein the data of the data sets is written in a binary format.

28. The method of claim 1, wherein the data of the data sets is written in a multi-state format.

29. The method of claim 1, wherein a portion of the data of the data sets is written in a binary format and a portion of the data of the data sets is written in a binary format.

30. The method of claim 1, wherein the data sets include includes categorical data and numerical data, and wherein said writing includes writing the categorical data in a binary format and numerical data in a multi-state format.

31. The method of claim 1, wherein each of the data sets are written into a plurality of NAND strings of a single corresponding bit line.

32. The method of claim 31, wherein the memory array is a flash memory array, wherein the plurality of NAND strings of the single corresponding bit lines are from multiple erase blocks.

33. The method of claim 1, wherein the memory array is formed on a first memory circuit including sensing circuitry connected to the bit lines, the sensing circuitry formed of a plurality of sensing units connectable to sense a first plurality of adjacent bits lines, wherein each data set is written along multiple bit lines from the first plurality of bit lines connectable to a corresponding sensing unit.

34. The method of claim 1, wherein the parameter values include “don’t care” values corresponding a sensing voltage for causing a memory cell to conduct for all data states.

35. The method of claim 1, further comprising:

writing a plurality of additional data sets selected by a host in a second memory array of a NAND architecture of NAND strings formed along bit lines and spanned by word lines, where the data sets are written oriented along the bit lines with corresponding data elements arranged to line along the same word lines, and wherein the additional data sets are arranged along bit lines of the second memory array so as to align with related data of the data sets as arranged on bit lines of the memory array;

concurrently biasing a plurality of word lines of the second array to a set of sensing voltage levels corresponding to host determined parameter values;

determining those of the bit lines of the second that conduct in response to said biasing of a plurality of word lines thereof to the set of sensing voltage levels; and

within data latches associated with the bit lines of the memory array, performing one or more logical/arithmetical operations upon the combined results of said determining for the memory array and the second array.

36. The method of claim 35, wherein the memory array and the second memory array are formed on different memory chips.

37. A method of operating a non-volatile array of a NAND architecture, the array having NAND strings formed along bit lines and spanned by word lines, where the bit lines span a plurality of the blocks and are each connected to a corresponding set of latches, the method comprising:

reading a first page of data from a first word line of the array;

storing the first page of data in a first of the latches for each of the bit lines of the array;

reading a second page of data from a second word line of the array, the first and second of the word lines being different;

storing the second page of data in a second of the latches for each of the bit lines of the array, wherein the first and second latches for each of the bit lines are different;

performing within said sets of latches one or more arithmetical/logical operations upon the contents of the first and second latches;

subsequently to performing said one or more arithmetical/logical operations, writing the result thereof from the sets of latches along a third word line of the array.

38. The method of claim 37, wherein said performing within said sets of latches one or more arithmetical/logical operations upon the contents of the first and second latches includes:

using a third latch for each of the bit lines of the array for carry operations, wherein the first, second and third latches for each of the bit lines are different.

39. The method of claim 37, further comprising:

writing the result of the one or more arithmetical/logical operations back to one or either the first and second latches for use in said writing result thereof.

40. The method of claim 37, wherein said one or more arithmetical/logical operations includes adding the contents of the first and second latches.

41. The method of claim 37, wherein said one or more arithmetical/logical operations includes subtracting the contents of the first and second latches.

42. The method of claim 37, wherein said one or more arithmetical/logical operations includes multiplying the contents of the first and second latches.

43. The method of claim 37, wherein said one or more arithmetical/logical operations includes dividing the contents of the first and second latches.

44. The method of claim 37, where the array is formed of multiple blocks of NAND strings formed along the bit lines, and wherein the third word line is from a different block than the first and second word lines.

45. The method of claim 44, wherein the first and second word lines are from different blocks.

46. The method of claim 44, wherein the first and second word lines are from the same block.

47. The method of claim 37, further comprising:
reading one or more additional pages of data from a corresponding word lines of the array; and

storing the one or more additional pages of data in latches for each of the bit lines of the array, wherein the one or more arithmetical/logical operations further involve the data of the one or more additional pages of data.

48. The method of claim 37, subsequently further comprising
reading a third page of data from the array;
storing the third page of data in the first of the latches for each of the bit lines of the array;

reading a fourth page of data from the array;
storing the fourth page of data in the second of the latches for each of the bit lines of the array;

subsequently performing within said sets of latches the one or more arithmetical/logical operations upon the contents of the first and second latches;

subsequently writing the result of the one or more arithmetical/logical operations from the sets of latches along a word line the array.

49. The method of claim 48, wherein the first, second, third and fourth pages of data are all read from different word lines of the array.

50. The method of claim 37, wherein the one or more arithmetical/logical operations include floating point operations.

51. The method of claim 37, further comprising:

prior to reading the first and second pages of data, writing a plurality of data sets selected by a host in the array, where the data sets are written oriented along the bit lines with corresponding data elements arranged to lie along the same word lines, the first and second pages of data being read from said data sets.

52. The method of claim 51, where the array is formed of multiple blocks of NAND strings formed along the bit lines, and wherein the first and second pages of are read respectively from first and second blocks, the second page of data begin read from a word line of the second block corresponding to a word line of the first block from which the first page of data was written.

53. The method of claim 52, wherein the second block holds an updated version of data sets in the first block.

54. The method of claim 52, wherein the second block holds data sets in the first block, but which are aligned along different bit lines than in the first block.




FIG. 1




FIG. 2





FIG. 3

4 / 22

	WL	BL BLB		BL BLB BL BLB			
		BL	BLB	BL	BLB	BL	BLB
Sense 1	Low	1	0	0	1	1	0
	Hi	1	0	0	1	0	1
	Hi	1	0	0	1	0	1
	Low	1	0	0	1	1	0
	Low	1	0	0	1	1	0
	Hi	1	0	0	1	0	1
	Low	1	0	0	1	1	0
Sense 2	sense 1	c	nc	nc	c	c	nc
	Hi	1	0	0	1	1	0
	Low	1	0	0	1	0	1
	Low	1	0	0	1	0	1
	Hi	1	0	0	1	1	0
	Hi	1	0	0	1	1	0
	Low	1	0	0	1	0	1
	Hi	1	0	0	1	1	0
Sense 2		c	nc	nc	c	nc	c
XOR1&2		1	1	1	1	0	0

Search for 00 pattern -- output column address --> data block


FIG. 4**FIG. 5**

5 / 22

	D0	D1	D3	D4	D5	D6	D7	D8	
Low	1	0			0	1	1	0	
Hi	1	0			0	1	0	1	
Hi	1	0			0	1	0	1	
Low	1	0			0	1	1	0	
Low	1	0			0	1	1	0	
Hi	1	0			0	1	0	1	
Low	1	0			0	1	1	0	
Hi	0	1			1	0	0	1	Reverse pattern
Low	0	1			1	0	1	0	
Low	0	1			1	0	1	0	
Hi	0	1			1	0	0	1	
Hi	0	1			1	0	0	1	
Low	0	1			1	0	1	0	
Hi	0	1			1	0	0	1	
Sense	nc	nc			nc	nc	c	nc	

FIG. 6A

6 / 22

**FIG. 6B**

7 / 22




FIG. 6C




FIG. 7

8 / 22




FIG. 8




FIG. 9




FIG. 10




FIG. 11








FIG. 12

12 / 22

**FIG. 13****FIG. 14**

13 / 22

**FIG. 15****FIG. 16**

14 / 22




FIG. 17

In range	XDL=0 UDL=0
Match	XDL=1 UDL=1
Out Bound	XDL=0 UDL=1

FIG. 18

Search for >010011001

Data	Read MSB9	Set 1	Read MSB8	Set 2	Read MSB6	Set 3	Final Search
100110111	1>0	XDL=1 UDL=1	ignore read	no update	no update		1
10100110	0=0	XDL=0 UDL=0	1=1	XDL=0 UDL=0	1>0	XDL=1 UDL=1	1
1011010	0=0	XDL=0 UDL=0	1>0	XDL=0 UDL=1	no update		0
11100	0=0	XDL=0 UDL=0	1>0	XDL=0 UDL=1	no update		0

FIG. 19

Search 123<x<231

	Read MSB1	Set 1	Read 2	Set 2	Read 3	Set 3	Read 3	Set 3	Final Search
467	4>2	XDL=0 UDL=1	ignore read	no update	no update		no update		0
246	2=2	XDL=0 UDL=0	4>3	XDL=0 UDL=1	no update		no update		0
132	1<2(upper)	XDL=0 UDL=0	1=1(lower)	XDL=0 UDL=0	3=3(upper)	XDL=1 UDL=1	3>2(upper)		1
34	0<2(upper)	XDL=0 UDL=0	0<1(lower)	XDL=0 UDL=1	no update		no update		0

FIG. 20




FIG. 21

FIG. 22




FIG. 23




FIG. 24

18 / 22

Example $10+3=13$

$$\begin{array}{r} 10 \\ 1010b \end{array} + \begin{array}{r} 3 \\ 0011b \end{array} = ?$$

		T0	T1	T2	T3
Block A	0	0			
	1		1		
	0			0	
	1				1

Block B	1	1			
	1		1		
	0			0	
	0				0

XDL "Carry" 0 1 0 0

Block C		1			
			0		
				1	
					1

1
0
1
1
=1101b=13

FIG. 25AExample $10-3=7$

$$\begin{array}{r} 10 \quad 1 \ 0 \ 1 \ 0 \\ 3c+1 \quad 1 \ 1 \ 0 \ 1 \\ \hline 7 \quad \boxed{1} \ 0 \ 1 \ 1 \ 1 \end{array} \quad \text{discard the overflow bits}$$

FIG. 25BThe multiplication can be done using bit shift and addition. 10×3

$$\begin{array}{r} 10 \quad 1 \ 0 \ 1 \ 0 \\ 3 \quad \times \quad 1 \ 1 \\ \hline \quad 1 \ 0 \ 1 \ 0 \\ 1 \ 0 \ 1 \ 0 \\ \hline 1 \ 1 \ 1 \ 1 \ 0 \end{array}$$

FIG. 25C

19 / 22




FIG. 26A




FIG. 26B




FIG. 27




FIG. 28




FIG. 29




FIG. 30

FIG. 31

21 / 22




FIG. 32




FIG. 33




FIG. 34