

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 August 2008 (21.08.2008)

PCT

(10) International Publication Number
WO 2008/100324 A9

(51) International Patent Classification:
H01L 25/065 (2006.01)

(21) International Application Number:
PCT/US2007/071370

(22) International Filing Date: 15 June 2007 (15.06.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/453,885 16 June 2006 (16.06.2006) US
11/538,567 4 October 2006 (04.10.2006) US

(71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, NY 10504 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **EMMA, Philip,**

George [US/US]; 28 Fox Den Road, Danbury, CT 06811 (US). **KNICKERBOCKER, John, Ulrich** [US/US]; 423 High Street, Monroe, NY 10950 (US). **PATEL, Chirag, S.** [US/US]; 201 Chateau Rive, Peekskill, NY 10566 (US).

(74) Agent: **WARDAS, Mark;** International Business Machines Corporation, T.J. Watson Research Center, Intellectual Property Law Department, 1101 Route 134 and Kitchawan Road (US).

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

[Continued on next page]

(54) Title: METHOD FOR ACHIEVING VERY HIGH BANDWIDTH BETWEEN THE LEVELS OF A CACHE HIERARCHY IN 3-DIMENSIONAL STRUCTURES, AND A 3- DIMENSIONAL STRUCTURE RESULTING THEREFROM




Fig. 2A

(57) Abstract: A method of electronic computing, and more specifically, a method of design of cache hierarchies in 3-dimensional chips, and a cache hierarchy resulting therefrom, including a physical arrangement of bits in cache hierarchies implemented in 3 dimensions such that the planar wiring required in the busses connecting the levels of the hierarchy is minimized. In this way, the data paths between the levels are primarily the vias themselves, which leads to very short, hence fast and low power busses.




Fig. 2B



(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

(48) Date of publication of this corrected version:

22 May 2009

(15) Information about Correction:

see Notice of 22 May 2009

**METHOD FOR ACHIEVING VERY HIGH BANDWIDTH BETWEEN THE LEVELS
OF A CACHE HIERARCHY IN 3-DIMENSIONAL STRUCTURES, AND A 3-
DIMENSIONAL STRUCTURE RESULTING THEREFROM**

5

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The present invention generally relates to a method for electronic computing, and more specifically, to a method for design of cache hierarchies in 3-dimensional chips, and 3-dimensional cache hierarchy structures resulting therefrom.

10

Description of the Related Art

[0002] The present invention provides a method in which a natural synergy is achieved by marrying two evolving fields of endeavor, which has not been recognized by the conventional methods.

15

[0003] First, recent work has demonstrated the viability of interconnecting two or more planes of circuits by thinning those planes (e.g., to a few hundred microns or less), etching dense via patterns in them, and then interconnecting them with metalization processes. The resulting structure is a monolithic “chip” including multiple planes of circuits. This advance is quite literally a new dimension in the scaling of circuit density.

20

[0004] Second, as circuit density has scaled, single chips have grown to contain more and more of the computer system. Two decades ago, it was a revelation that an entire processor could fit on a single chip. At the 180 nanometer CMOS node, it was a revelation that not only the processor’s Level-1 cache (L1) was contained, but for the first time it was also feasible to

include the next level of cache, L2, on the chip with the processor. Additionally, about a decade ago, the first single-chip multiprocessors began being produced.

[0005] At densities facilitated at the 90 nanometer node and beyond, together with the aforementioned ability to create 3-dimensional structures, single chip systems of the future will contain not only multiple processors, but also multiple levels of the cache hierarchy.

[0006] The access time of a cache is determined, to a large extent, by its area. Therefore, a processor's Level-1 cache (L1), which is integral to the processor pipeline itself is kept small so that its access time is commensurate with the processor speed, which today can be several Gigahertz. Because the L1 is small, it cannot contain all of the data that will be used by the processor when running programs. When the processor references a datum that is not contained within the L1, this is called an L1 "cache miss."

[0007] In the event of an L1 miss, the reference is forwarded to the next level in the hierarchy (say, L2), to determine whether the datum is there.

[0008] If the requested datum is in the L2 cache, then data (including the datum that was specifically referenced) are moved from the L2 cache to the L1 cache, and the original reference is satisfied.

[0009] If the referenced datum is not in the L2 cache, then this is also a "cache miss" (an L2 "cache miss"), and the reference continues to percolate up the hierarchy (e.g., say, to L3 and above). By convention, higher levels in a cache hierarchy are physically larger (and hence, hold more data), and thus, progressively slower.

[0010] Data in a cache are stored in "lines," which are contiguous chunks of data (i.e., being a power-of-2 number of bytes long, aligned on boundaries corresponding to this size). Thus, when a cache miss is serviced, it is not merely the specific datum that was requested that is

moved down the hierarchy. Instead, the entire cache line containing the requested datum is moved. Data are stored in “lines” for two reasons.

[0011] First, each entry in a cache has a corresponding directory entry (e.g., in a cache directory) that contains information about the cache entry. If each byte in a cache were given its own entry, there would be a prohibitive number of directory entries (i.e., equal to the number of bytes in the cache) making the administrative overhead for the cache (the directory) huge. Thus, instead, there is one directory entry per cache line (which is typically between 32-256 bytes today, depending on the processor).

[0012] Second, program reference patterns exhibit what is called “spatial locality of reference,” meaning that if a particular datum is referenced, then it is very likely that other data that are physically proximate (e.g., by address) to the referenced datum will also be referenced. Thus, by bringing in an entire cache line, more of the spatial context of the program is captured, thereby reducing the number of misses.

[0013] The bandwidth between levels in a hierarchy is equal to the amount of data that is moved per unit of time. It is noted that the bandwidth includes both necessary movement (e.g., the data that are actually used) and unnecessary movement (e.g., data that are not ever referenced).

[0014] To achieve high performance in a processor, it is important to not take many misses, since misses are a dominant component of delay. One method of reducing the number of misses incurred is to anticipate what might be used by a running program, and to “prefetch” that data down the hierarchy before it is referenced. In this way, if the program actually does reference what was anticipated, there is no miss. However, the more a processor speculates about what might be referenced (so as to prefetch), the more unnecessary movement takes place, since some

of what is anticipated will be wrong. This means that facilitating higher performance by the elimination of misses will require more bandwidth.

[0015] The actual bandwidth used (as defined above) is the amount of data moved per unit time. If the program runs fast, the unit of time will be shorter, hence the bandwidth higher.

5 Notice the distinction between the actual bandwidth, and the “bandwidth capacity,” which is equal to the maximum amount of data that could be moved if the busses were 100% utilized. Bandwidth capacity is equal to the width of the bus (in bytes) times the bus frequency.

[0016] For example, if an 8-byte bus runs at 1 Gigahertz, then the bandwidth capacity of the bus is 8 Gigabytes per second. It is noted also that if the processor runs at 2 Gigahertz, then the
10 bus is 2 times slower than the processor. And if the cache line size is 128 bytes, then the 8-byte bus requires 16 bus cycles (which is $2 \times 16 = 32$ processor cycles) to move a cache line. Some of the data that is moved during these 32 processor cycles is useless. Further, if a subsequent miss occurs during the large window in which this cache line is being moved, the subsequent miss can be further delayed by the bus transfer that is already in progress.

15 **[0017]** For this reason, it is important to have a bandwidth capacity that is much larger than the actual bandwidth demand. Very high bandwidth facilitates two things that are crucial to high performance computing.

[0018] First, very high bandwidth allows cache lines to be transferred very quickly so that the transfers will not interfere with other miss traffic in the system. (For example, if the bus
20 above were 128 bytes wide and ran at 3 Gigahertz, it could transfer the cache line in a single processor cycle.)

[0019] Second, having an ample surplus of bandwidth capacity facilitates operations like prefetching, which will place a much higher bandwidth demand on the system.

[0020] The reason that bus widths tend to be much narrower than cache line sizes (e.g., 8 bytes instead of 128 bytes) is that planar wiring capability is limited. Because these busses tend to be long, they also tend to be wired in high-level (e.g., relatively thick or fat) wire to minimize resistance so as to maximize speed. Wide busses (e.g., much wider than 8 bytes) would impose considerable blockages on the upper levels of metal, so they are generally not used.

[0021] Busses tend to be slower than processors (e.g., 1 Gigahertz instead of 2 Gigahertz) because they are too long (e.g., 5-10 millimeters or more) since they are connecting large aerial structures (caches) in a plane.

SUMMARY OF THE INVENTION

[0022] In view of the foregoing and other exemplary problems, drawbacks, and disadvantages of the conventional methods and structures, an exemplary feature of the present invention is to provide a method which maximize the number of interconnections between the levels in a cache hierarchy so as to maximize the bus widths, and the cache hierarchy resulting from such a novel method.

[0023] Particularly, the present invention provides a method in which a natural synergy is achieved by marrying two evolving fields of endeavor, which has not been recognized by the conventional methods.

[0024] An exemplary feature of the present invention is to minimize the bus wire lengths between the levels in a cache hierarchy so as to maximize the bus speed and so as to minimize the energy required per transfer.

[0025] In a first aspect of the present invention, a method for arranging bits within a cache hierarchy implemented on multiple physical planes such that horizontal wiring distances in intra-

level busses are minimized, includes physically partitioning each cache level into cache islands, each cache island including a subset of congruence classes, wherein the partitioning is performed in correspondence across cache levels such that a "cache island" containing specific congruence classes at one cache level are directly above or below the corresponding cache island which
5 contains the corresponding congruence classes of an adjacent cache level, i.e., physically positioning each cache island directly over the corresponding cache islands of different cache levels.

[0026] In yet another aspect of the present invention, a method of deploying computing infrastructure in which recordable, computer-readable code is integrated into a computing
10 system, and combines with the computing system for arranging bits within a cache hierarchy implemented on multiple physical planes such that horizontal wiring distances in intra-level busses are minimized, includes physically partitioning each cache level into cache islands, each cache island including a subset of congruence classes, wherein the partitioning is performed in correspondence across cache levels such that the congruence classes within a cache island at one
15 cache level correspond to a same congruence classes within a corresponding cache island at another cache level, and physically positioning each cache island directly over the corresponding cache islands of different cache levels.

[0027] In another aspect of the present invention, a computer-readable medium tangibly embodying a program of machine readable instructions executable by a digital processing
20 apparatus performs a method program for causing a computer to perform the exemplary methods described herein.

[0028] In the present invention, a method is provided for the physical arrangement of the bits in cache hierarchies implemented in three (3) dimensions, so as to minimize the planar wiring

required in the busses connecting the levels of the hierarchy. In this way, the data paths between the levels are primarily determined by the vias themselves. This leads to very short, (and hence fast) and low power busses.

[0029] In the present invention, the bits within each level of a cache hierarchy are so arranged such that transfers between levels in the hierarchy do not involve very much of the planar wiring in any level. The present invention restricts bit placement within the plane by grouping the bits in accordance with two (orthogonal) criteria such that the movement between cache levels is nearly completely vertical in a 3-dimensional structure.

[0030] For example, the present inventors have recognized that an emerging technology is the ability to make 3-dimensional structures by thinning chips, laminating them together, and interconnecting them with through-vias in an areal grid. However, such structures provide very wide interconnecting busses.

[0031] The present invention provides for interconnecting multiple planes of memory together on a wide vertical bus to provide unprecedented amounts of bandwidth between adjacent levels in a cache hierarchy. The present invention provides a method for organizing the bits within the arrays of the adjacent levels so that minimal x-y wiring is required to transfer data. In this way, the wide bus can be almost entirely vertical, and also can be very short (e.g., a few hundreds microns). This provides a bus that can run at much higher speed, and much lower power than conventional 2-dimensional busses.

BRIEF DESCRIPTION OF THE DRAWINGS

[0032] The foregoing and other exemplary purposes, aspects and advantages will be better understood from the following detailed description of exemplary aspects of the invention with reference to the drawings, in which:

5 [0033] Figure 1 exemplarily illustrates a cache hierarchy showing an L1.5 SRAM plane and an L2 including two planes of eDRAM;

[0034] Figure 2 logically illustrates cache geometries for L1.5 and L2 caches, according to exemplary aspects of the present invention;

10 [0035] Figure 3 exemplarily illustrates a bitline structure in which 32 bits in each of the eDRAM planes can drive the corresponding 8 bits on the SRAM plane;

[0036] Figures 4A-4B exemplarily illustrate a multicore example of 8 processors, each with a private cache hierarchy and some intra-processor interconnection;

[0037] Figure 5 illustrates an exemplarily method according to the present invention; and

[0038] Figure 6 illustrates another exemplarily method according to the present invention.

15

DETAILED DESCRIPTION OF EXEMPLARY

ASPECTS OF THE INVENTION

[0039] Referring now to the drawings, and more particularly, to Figures 1-6, there are shown exemplary aspects of the method and structures according to the present invention.

20 [0040] For the purposes of illustrating the principle concepts of the present invention, herein will be first focused on a simple 2-level hierarchy comprising an L1.5 cache and an L2 cache. Roughly speaking, it is a good rule of thumb that adjacent levels in a hierarchy have a capacity

difference of an order of magnitude or so. In this example, the L1.5 is taken to be 1 Megabyte, and the L2 is 8 Megabytes.

[0041] First, these caches can be arranged so that they are in “aerial correspondence.” That is, the areas of the two caches can be made approximately the same. One of them can be placed on top of the other. This is preferred (e.g., required) so that when moving data between the two levels, large horizontal movements are not required. Preferably, substantially only vertical movements will be needed. Then, the bits of the two caches will be arranged so that the physically corresponding bits (between the two levels) are in identical x-y (planar) proximities. This will eliminate the need for much x-y (planar) wiring in the intra-planar communication paths.

[0042] Although the capacity of the L2 is (in this case) 8X larger than the capacity of the L1.5, the invention preferably makes them aerially commensurate by using different circuit densities, and by using multiple planes of storage.

[0043] For example, let the L1.5 be implemented in 6T SRAM, and let the L2 be implemented in embedded DRAM (eDRAM). For argument’s sake, assume that eDRAM offers a 4X density improvement over SRAM. To facilitate an 8X capacity difference then requires two planes of eDRAM to implement the L2 to correspond to a single plane of SRAM for the L1.5. Figure 1 shows this basic arrangement.

[0044] That is, Figure 1 illustrates a cache hierarchy structure 100 showing an L1.5 SRAM plane 110 (having for example 1 MB of SRAM) and an L2 cache 120 comprising two planes of eDRAM 120A, 120B (having 4 MBs for example).

[0045] Assuming a 0.65 square micron SRAM cell with a 75% array efficiency, 1 Megabyte of SRAM is approximately 8.2 square millimeters. Using this number, a total area in the

neighborhood of 10 square millimeters can be aimed for (after having de-rated the array efficiency to account for rearranging the bits, and after having put in an aerial array of vias).

[0046] In this exemplary embodiment, a line size of 256 bytes is chosen. The 1 Megabyte L1.5 then has 4K lines, and the 8 Megabyte L2 has 32K lines. The set associativity of the L1.5 is
5 chosen as 8-way, and the set associativity of the L2 is chosen as 16-way. This makes the number of congruence classes in the L1.5 be $4K / 8 = 512$, and the number of congruence classes in the L2 be $32K / 16 = 2K$. These geometries are shown in Figure 2.

[0047] That is, Figure 2 logically illustrates exemplary cache geometries 210, 220 for L1.5 and 2 caches, respectively.

10 [0048] Because this is a composite L1.5/L2 structure, a combined L1.5/L2 directory is chosen to be constructed, in SRAM, and put it on the SRAM plane, although this is arbitrary, and not essential to the present invention. Basically, a “combined” directory is really an L2 directory (having 32K entries), where each entry has an additional field to specify its status (MESI) in the L1.5, in the case that the corresponding data is resident in the L1.5. This was estimated to take
15 about 1.6 square millimeters, assuming 8 bytes for a directory entry.

[0049] On a directory search, if there is an L1.5 miss and an L2 hit, then the directory should identify the line in the L2 that is to be brought into the L1.5, and it should choose the location (set within congruence class) in the L1.5 into which it is to be placed. Therefore, it should select one of 32K lines in the L2, which it can do by transmitting as few as 15 bits up to the L2 plane,
20 or as many as 32K select signals, depending on how many vias are desired for use, and where one wishes to place decoders. The choice here is arbitrary (with regard to teaching this invention), albeit important to the goals of the design. The point is that the number of vias used for this can be small.

[0050] Initially, a goal is to maximize the width of the data bus. For a 256-byte line, were a line to be moved in a single transfer, it would require 2304 signals (256 bytes X 9 bits per byte). Using one signal via per power via requires a total count of about 5000 vias to do this (together with address selects and various control signals). Using a via size of 16 square microns yields a
5 total via area of 0.08 square millimeters. So in fact, if allowed a total via area overhead of 20% (of the 10 square millimeter area target), 32 such busses could be placed in use, and as many as 32 cache lines (this is 8 Kilobytes) to be moved simultaneously in an area of about 12 square millimeters.

[0051] While this amount of bandwidth capacity is staggering by today's standards, the
10 question remains as to how to route this many signals (say $2500 \times 32 = 80K$) in a feasible way, and so that those wires are very short.

[0052] The first observation made is that, of all of the bits in the L2 (8 Megabytes X 9 bits per byte = 72 Megabits) and of all of the bits in the L1.5 (9 Megabits), it is not the case that any of the 72 Megabits in the L2 can be moved to any of the 9 Megabits in the L1.5. In fact, for any
15 particular bit in the L1.5, there are only 64 bits in the L2 that can map into that particular bit. This is true for two independent reasons as described below. The present invention physically arranges the bits of the L1.5 and L2 so that only those bits of the L2 that can map to a bit in the L1.5 are above it.

[0053] A first (and trivial) reason for this mapping limitation, is that each bit in a cache line
20 occupies a unique position. That is, a line is a contiguous and ordered group of 256 bytes, and each of those bytes is a contiguous and ordered group of 9 bits. Therefore, there are 2304 unique bit positions in a cache line. When a line is moved from the L2 to the L1.5, bit 0 of the L2 line is placed into bit position 0 in the L1.5 line, bit 1 of the L2 line is placed into bit position 1 in the

L1.5 line, and so on. Bit 0 (and all other bits) can only be placed into its corresponding bit position. It cannot be placed into any other bit position.

[0054] Secondly, and less obviously, for any congruence class in the L1.5, there are only a handful of congruence classes in the L2 that can map to it. Specifically, if there are n_1 congruence classes in the L1.5, and n_2 congruence classes in the L2, then only n_2/n_1 L2 congruence classes can map to a given L1.5 congruence class. In the example being shown, the numbers are $2K / 512 = 4$ L2 congruence classes per L1.5 congruence class. Further, those L2 congruence classes cannot map to any other L1.5 congruence class.

[0055] To put it differently, the L1.5/L2 combination can be viewed as 512 (the number of L1.5 congruence classes) independent L1.5/L2 sub-caches, each in its own little island with its own bus between the L1.2 and L2. Each subcache is an L1.5 congruence class and the set of L2 congruence classes that map to it. Therefore, the 1M L1.5 / 8M L2 aggregation can be thought of as 512 independent 2K L1.5 / 16K L2 aggregations, and each of those have 2304 independent bit positions.

[0056] It is noted that a line coming from the L2 (which can reside in any of the 16 sets of the 4 congruence classes of its subcache) can be placed into any of the 8 sets of the L1.5 congruence class. If the cache is partitioned down to the bit position, it means that there are $16 \times 4 = 64$ bits in the L2 that can map to any of 8 bits (the set associativity of the L1.5) in the L1.5. Thus, conceivably, the smallest "island" in this cache would correspond to an 8-bit island of the L1.5 residing underneath 2 (recall that there are 2 eDRAM planes) 32-bit islands of the eDRAM.

[0057] This situation is shown in Figure 3. That is, Figure 3 exemplarily illustrates a bitline structure 300 in which 32 bits in each of the eDRAM planes 310, 320 can drive the corresponding 8 bits on the SRAM plane 330.

[0058] In Figure 3, the data path is shown as having independent signals for each eDRAM plane. For a fixed number of vias, this would imply half the bus width (per line) than what was described above, although in this configuration, two lines could be moved at the same time (e.g., one from each eDRAM plane). Alternatively, a single via could be used, and only one plane
5 could be selected at a time for transmission.

[0059] While these particular partitionings are too small to lead to good array efficiencies, they demonstrate that the total x-y (planar) motion required for an L2-L1.5 transfer is very small indeed, i.e., it can be made almost completely vertical. It is noted that since the vertical distance is small, e.g., 400 microns for the three 200-micron thick planes shown here, the transfer speeds
10 can be much faster than what would be achievable in several millimeters of horizontal wiring (i.e., than what is possible in a plane). Additionally, since the capacitance is very small, the energy required per transfer is also much smaller.

[0060] Assuming that a real design would tradeoff array efficiency with the size of the islands (which determines the length of horizontal wire required in a transfer), the table below
15 shows several ways to partition the caches, and the busses that result using 80K signal vias:

# Congruence Classes Contained per Island	Bus Width per Island	# Bus Cycles to Transfer a Line	Total # of Lines in Flight Simultaneously
1	16	16	512
2	32	8	256
4	64	4	128
8	128	2	64
16	256	1	32

[0061] The last row in the table basically represents the situation implied above - using 80K signal vias. That is, for a 256-byte bus width, we can transfer a 256-byte line in a single cycle.

20 With our 80K signal vias, we can have 32 such busses, hence can have 32 lines in flight

simultaneously. Partitioning the 512 congruence classes into 32 islands (each with its own bus) would put 16 congruence classes into each island - as shown in the first column.

[0062] On the other end of the spectrum, the first row shows that if one partitioned the cache into 512 islands, each having a single congruence class, then with 80K signal vias, each of the islands could only have a 16 byte bus, and it would take 16 cycles to move a 256 byte line. In this case, 512 lines can be in flight at the same time, but it takes longer to move a line.

[0063] Each row in the table shows a different partitioning, but each of the partitionings has the same bandwidth capacity: 8 Kilobytes per cycle.

[0064] While the above descriptions pertain to a single L1.5 / L2 pair, the question arises as to how this same reasoning applies to a multicore system. Figures 4A-4B illustrate a multicore example of 8 processors, each with a private hierarchy and some intra-processor interconnection.

[0065] That is, Figures 4A-4B show a generalization of a processor chip in which there are 8 processors, each with a private L1.5 cache. Figure 4A shows the 8 processors with the L1.5s. Of course, the area is dominated by the L1.5s (which is not apparent in the conceptual view shown).

[0066] While any number of processors is acceptable, the reason that 8 was chosen in this example is that the chip can be partitioned into 9 similar regions, with the outside 8 regions each holding a processor and its cache, where the central region can be used for a switch, to handle traffic between the caches.

[0067] Figure 4B shows that same 8-processor chip on the bottom of a stack of 4 memory chips (e.g., eDRAM) 440A-440D, although any number is feasible. Just as the bottom processor chip is partitioned into 9 physical regions, each of the memory chips above is so partitioned as well (although this isn't explicitly shown). Essentially, this system is just 8 copies of the single

system described above, with some switching (at least on the processor chip, but perhaps on the memory chips as well) to facilitate intra-processor communications.

[0068] Figure 5 shows an exemplarily method 500 according to the present invention for arranging bits within a cache hierarchy implemented on multiple physical planes such that horizontal wiring distances in intra-level busses are minimized. As shown in Figure 5, the method can include co-designing physical structures of cache levels for optimizing interconnections between logically adjacent levels of the cache hierarchy, wherein the cache levels are physically positioned over each other.

[0069] Figure 6 shows an exemplarily method 60 according to the present invention for arranging bits within a cache hierarchy implemented on multiple physical planes such that horizontal wiring distances in intra-level busses are minimized. As shown in Figure 6, the method can include physically partitioning each cache level into cache islands (e.g., step 610). It is noted that each cache island preferably can include a subset of congruence classes. It is also noted that the partitioning is performed in correspondence across cache levels such that the congruence classes within a cache island at one cache level map to same congruence classes of a corresponding cache island at a different cache level. The method 600 includes physically positioning each cache island directly over the corresponding cache islands of different cache levels (e.g., step 620).

[0070] While the invention has been described in terms of several exemplary aspects, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

[0071] Further, it is noted that, Applicant's intent is to encompass equivalents of all claim elements, even if amended later during prosecution.

CLAIMS**What is claimed is:**

- 1 1. A method of arranging bits within a cache hierarchy implemented on multiple physical
2 planes such that horizontal wiring distances in intra-level busses are minimized, said method
3 comprising:
4 co-designing physical structures of cache levels for optimizing interconnections between
5 logically adjacent levels of said cache hierarchy,
6 wherein said cache levels are positioned over each other.
- 1 2. The method according to claim 1, wherein said co-designing comprises:
2 physically partitioning each said cache level into cache islands, each said cache island
3 including a subset of congruence classes,
4 wherein said partitioning is performed in correspondence across cache levels such
5 that the congruence classes within a cache island at one cache level map to same congruence
6 classes of a corresponding cache island at a different cache level.
- 1 3. The method according to claim 1, wherein said bits within said cache hierarchy are
2 arranged to minimize horizontal transfers between said cache levels in the cache hierarchy.
- 1 4. The method according to claim 1, wherein adjacent cache levels are arranged to include
2 different capacities.

1 5. The method according to claim 1, wherein adjacent cache levels are arranged to include
2 same capacities.

1 6. The method according to claim 1, wherein one of said cache levels in the cache hierarchy
2 comprises a plurality of physical planes.

1 7. The method according to claim 1, wherein areas of each cache level are arranged to be
2 substantially the same.

1 8. The method according to claim 2, wherein said partitioning is selected to optimize a size
2 of each said cache island.

1 9. A method of deploying computing infrastructure in which recordable, computer-readable
2 code is integrated into a computing system, and combines with said computing system to
3 perform the method according to Claim 1.

1 10. A computer-readable medium tangibly embodying a program of machine readable
2 instructions executable by a digital processing apparatus to perform a method of causing a
3 computer to perform the method according to Claim 1.

1 11. A method of arranging bits within a cache hierarchy implemented on multiple physical
2 planes such that horizontal wiring distances in intra-level busses are minimized, said method
3 comprising:

physically partitioning each cache level into cache islands, each said cache island including a subset of congruence classes, wherein said partitioning is performed in correspondence across cache levels such that the congruence classes within a cache island at one cache level map to same congruence classes of a corresponding cache island at a different cache level; and positioning each said cache island over the corresponding cache islands of different cache levels.

12. The method according to claim 11, wherein said bits within said cache hierarchy are arranged to minimize horizontal transfers between said cache levels in the cache hierarchy.

13. The method according to claim 11, wherein adjacent cache levels are arranged to include different capacities.

14. The method according to claim 11, wherein adjacent cache levels are arranged to include same capacities.

15. The method according to claim 11, wherein one of said cache levels in the cache hierarchy comprises a plurality of physical planes.

16. The method according to claim 11, wherein areas of each cache level are arranged to be substantially the same.

1 17. The method according to claim 11, wherein said partitioning is selected to optimize a size
2 of each said cache island.

1 18. A method of deploying computing infrastructure in which recordable, computer-readable
2 code is integrated into a computing system, and combines with said computing system to
3 perform the method according to Claim 11.

1 19. A computer-readable medium tangibly embodying a program of machine readable
2 instructions executable by a digital processing apparatus to perform a method of causing a
3 computer to perform the method according to Claim 11.

1 20. A cache hierarchy implemented on multiple physical planes such that horizontal wiring
2 distances in intra-level busses are minimized, said cache hierarchy comprising:

3 a plurality of cache levels physically partitioned into cache islands, each said cache island
4 including a subset of congruence classes,

5 wherein said cache levels are partitioned in correspondence across cache levels such that
6 the congruence classes within a cache island at one cache level map to same congruence classes
7 of a corresponding cache island at a different cache level, and

8 wherein said cache islands are positioned over the corresponding cache islands of
9 different cache levels.

1 21. The cache hierarchy according to claim 20, wherein adjacent cache levels include
2 different capacities.

1 22. The cache hierarchy according to claim 20, wherein adjacent cache levels include same
2 capacities.

1 23. The cache hierarchy according to claim 20, wherein areas of each cache level are
2 substantially the same.

1 24. The cache hierarchy according to claim 20, wherein one of said cache levels in the cache
2 hierarchy comprises a plurality of physical planes.

1 25. The cache hierarchy according to claim 20, wherein said partitioning is selected to
2 optimize a size of each said cache island.

1 26. The cache hierarchy according to claim 20, wherein said cache hierarchy is arranged to
2 minimize horizontal transfers between said cache levels in the cache hierarchy.

1 27. The cache hierarchy according to claim 20, wherein data paths between said cache levels
2 comprise vias.

1 28. A design tool for designing a cache hierarchy according to claim 20, which is
2 implemented on multiple physical planes such that horizontal wiring distances in intra-level
3 busses are minimized.

1 29. A design tool comprising:

a cache hierarchy according to claim 20, which is implemented on multiple physical planes such that horizontal wiring distances in intra-level busses are minimized.

30. A computer system, comprising:

a cache hierarchy implemented on multiple physical planes such that horizontal wiring distances in intra-level busses are minimized, said cache hierarchy comprising:

a plurality of cache levels physically partitioned into cache islands, each said cache island including a subset of congruence classes,

wherein said cache levels are partitioned in correspondence across cache levels such that the congruence classes within a cache island at one cache level map to same congruence classes of a corresponding cache island at a different cache level, and

wherein said cache islands are positioned over the corresponding cache islands of different cache levels.

31. A computer system according to claim 30, wherein said cache islands are physically positioned directly over the corresponding cache islands of different cache levels.




Fig. 1




Fig. 2A




Fig. 2B

300




Fig. 3




Fig. 4A


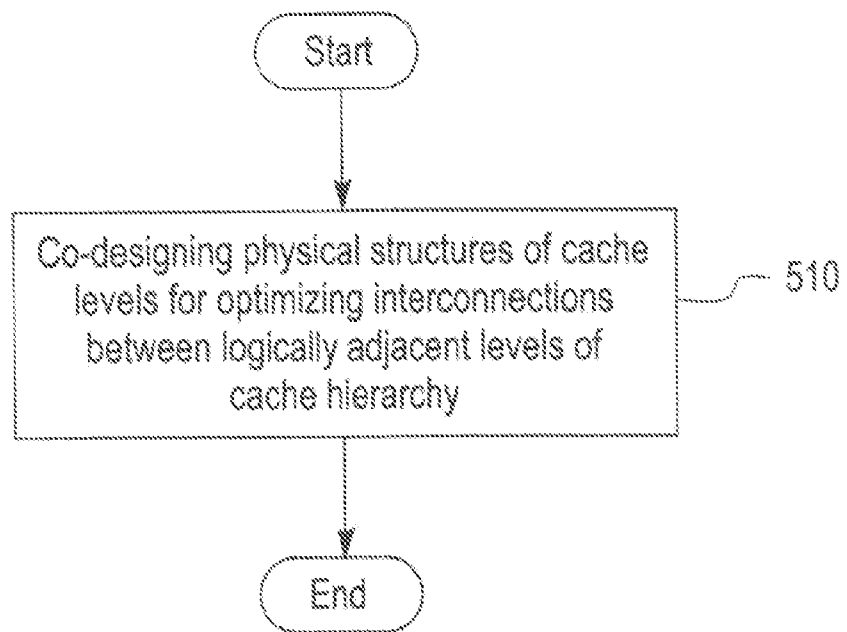
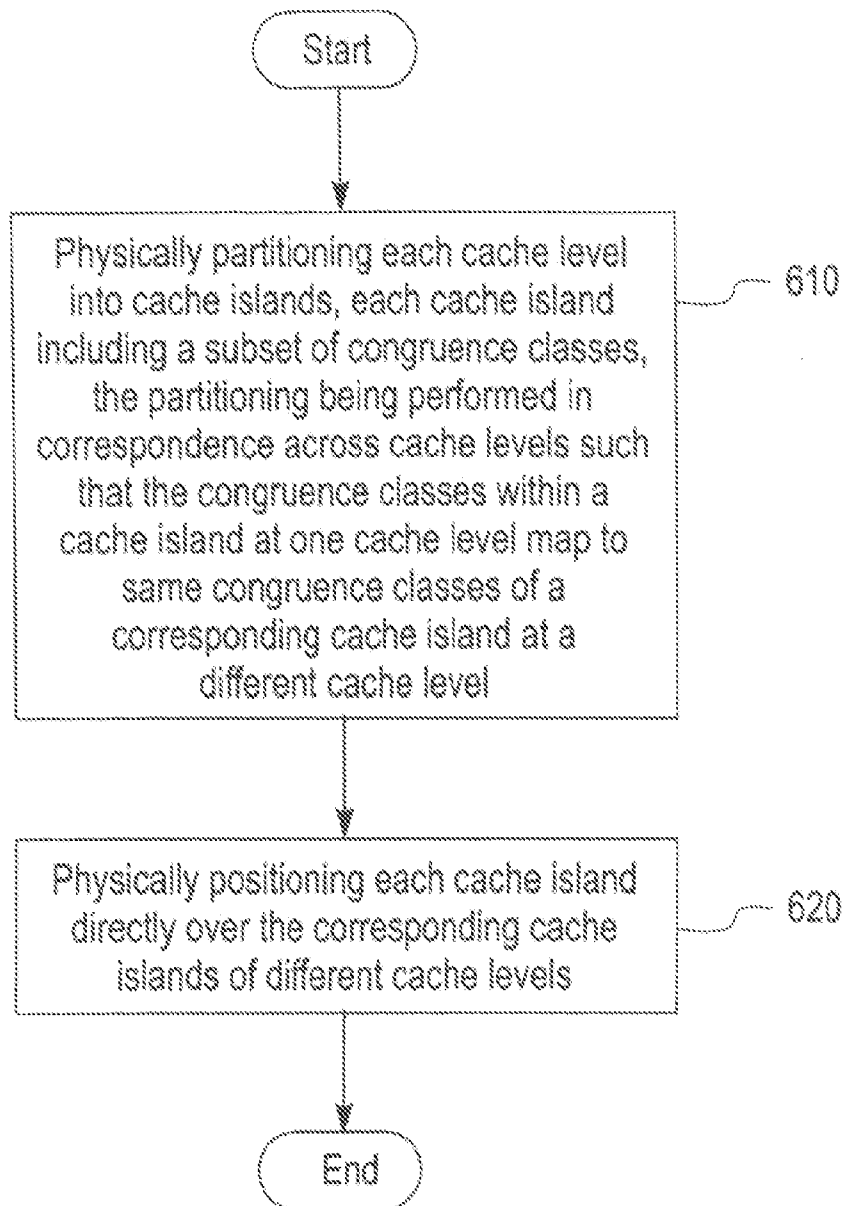


Fig. 4B

500**Fig. 5**

600**Fig. 6**