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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
30 November 2017 (30.11.2017)

WIPO | PCT

(10) International Publication Number
WO 2017/204928 A1

(51) International Patent Classification:
G06F 1/32 (2006.01)

(21) International Application Number:
PCT/US2017/027968

(22) International Filing Date:
17 April 2017 (17.04.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/162,452 23 May 2016 (23.05.2016) US

(71) Applicant: **QUALCOMM INCORPORATED** [US/US]; ATTN: International IP Administration, 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

(72) Inventors: **SAHU, Satyanarayana**; 5775 Morehouse, San Diego, California 92121 (US). **RAJ, Satish**; 5775 Morehouse Drive, San Diego, California 92121 (US). **CHANDRASEKARAN, Shiva Ram**; 5775 Morehouse Drive, San Diego, California 92121 (US). **QIU, Li**; 5775 Morehouse, San Diego, California 92121 (US). **TYAGI, Arun**; 5775 Morehouse Drive, San Diego, California 92121 (US). **PHILIP, Mathew**; 5775 Morehouse Drive, San Diego, California 92121-1714 (US). **VERMA, Rajesh**; 5775 Morehouse Drive, San Diego, California 92121 (US).

(74) Agent: **KELTON, Thomas W.** et al.; Haynes and Boone, LLP, 2323 Victory Avenue, Suite 700, Dallas, Texas 75219 (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, DJ, 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, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, 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,

WO 2017/204928 A1

(54) Title: SYSTEMS AND METHODS TO SEPARATE POWER DOMAINS IN A PROCESSING DEVICE

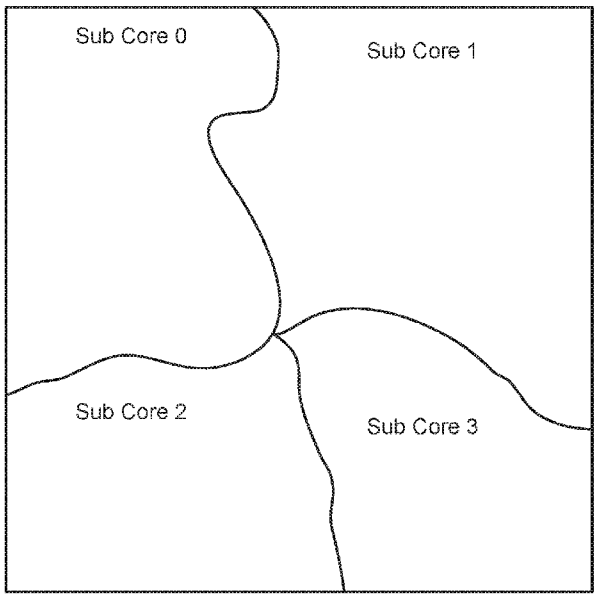

FIG. 1

(57) Abstract: A semiconductor device includes: a processing core having a plurality of sub cores, a plurality of power rails spanning from a first sub core to a second sub core of the plurality of sub cores, the plurality of power rails configured to provide an operating voltage to each of the first sub core and the second sub core, and a plurality of cells defining a boundary between the first sub core and the second sub core, each of the cells providing a discontinuity in a respective power rail, wherein the discontinuity includes a break in the respective power rail in more than one layer of the semiconductor device.

[Continued on next page]

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, ST, 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).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

— *with international search report (Art. 21(3))*

SYSTEMS AND METHODS TO SEPARATE POWER DOMAINS IN A PROCESSING DEVICE

Satyanarayana Sahu, Satish Raj, Shiva Ram Chandrasekaran, Li Qiu, Arun Tyagi, Mathew Philip, Rajesh Verma

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application claims priority to U.S. Nonprovisional Application No. 15/162,452, filed May 23, 2016, which is hereby incorporated by reference in its entirety as if fully set forth below in its entirety and for all applicable purposes.

TECHNICAL FIELD

[0002] This application relates to processing device design and, more specifically, to separating sub cores and power domains in a processing device.

BACKGROUND

[0003] A mobile computing device, such as a smart phone, contains a multi-core chip to provide computing power. Examples of processing cores include a Digital Signal Processor (DSP) core, a Graphics Processing Unit (GPU), a Central Processing Unit (CPU), a modem, and a camera core. Each core may have multiple sub cores. For instance, a GPU may include shaders and accumulators as sub cores, a DSP and a CPU may each include different processing units as sub cores.

[0004] Furthermore, each of the different sub cores may belong to a separate power domain. As an illustration, a core may be programmed to save power consumption by reducing a voltage to a sub core or turning power off altogether for a sub core that is not in use at a given time. When that sub core is in use, the processor restores full power to that sub core. In other words, each of the different sub cores are powered independently and may be turned on and off independently.

[0005] In one conventional example, each of the different sub cores is shaped as a rectangle or rectilinear shape. During design of the core, each of the different sub cores belongs to a different team, and each team is assigned a particular shape to work with so that the different sub cores may be placed together. Each team designs its sub core to fit the shape it is given, with the goal that the shapes should fit together in the final product. These well-defined rectangles or rectilinear shapes, separated from each other by some amount of area, simplify power domain partitioning. However, the space between the shapes may sometimes be considered wasted, whereas space on a semiconductor die is valued highly as miniaturization of devices is a priority in some applications.

[0006] Furthermore, the design process including assigning design teams to particular shapes can be well compartmentalized and efficient, but it may require multiple iterations when one or more of the design teams determines that the shape it is given may not allow it to meet constraints. Constraints include, e.g., data timing constraints and power constraints. The iterations may require negotiations among the different design teams and a manager of the core project. After several iterations, the design may be ready to be assembled together and to be taped out.

[0007] There is currently a need for a design that is more efficient with respect to space between cores and for a design process that is more efficient and has fewer iterations.

SUMMARY

[0008] Various embodiments include systems and methods that implement sub cores using irregular shapes on a processing device. Such irregular shapes may be facilitated by the use of power breaker cells, which provide discontinuities within power rails so as to separate and different power domains.

[0009] In one embodiment, a semiconductor device includes: a processing core having a plurality of sub cores, a plurality of power rails spanning from a first sub core to a second sub core of the plurality of sub cores, the plurality of power rails configured to provide an operating voltage to each of the first sub core and the second sub core, and a plurality of cells defining a boundary between the first sub core and the second sub core, each of the cells providing a discontinuity in a respective power rail, wherein the discontinuity includes a break in the respective power rail in more than one layer of the semiconductor device.

[0010] In another embodiment, a semiconductor chip includes: a plurality of processing cores disposed within semiconductor material of the semiconductor chip, a first one of the cores having a first sub core and a second sub core, a power rail spanning from the first sub core to the second sub core, the power rail including conductive lines within the semiconductor material, and a cell abutting the first sub core and the second sub core on the power rail, the cell providing a discontinuity in the power rail at a plurality of metal layers of the conductive lines.

[0011] In another embodiment, a semiconductor chip has a processing core, the processing core including a first sub core, a second sub core, means for distributing power to the first sub core and the second sub core, and means for providing discontinuities in the power distributing means at multiple metal layers of the semiconductor chip, wherein a boundary between the first sub core and the second sub core is defined by the discontinuity providing means abutting the first core and the second core.

[0012] In yet another embodiment, a method includes: power collapsing a first sub core while providing power to a second sub core, wherein the first sub core and the second sub core are implemented in a processing core on a semiconductor chip, wherein

a boundary between the first sub core and the second sub core includes a plurality of cells providing discontinuities in a plurality of power rails at multiple layers of the semiconductor chip; and power collapsing the second sub core independently of power collapsing the first sub core.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] Figure 1 is a simplified diagram illustrating an example processing core having multiple sub cores, according to one embodiment.

[0014] Figure 2 is a simplified diagram of a power breaker cell adapted according to one embodiment.

[0015] Figure 3 is an illustration of an example power breaker cell adapted according to one embodiment.

[0016] Figure 4 is an illustration of an example power breaker cell adapted according to one embodiment.

[0017] Figure 5 shows an example implementation, where a power breaker cell abuts other cells at a power domain boundary, according to one embodiment.

[0018] Figures 6 and 7 illustrate example implementations, wherein multiple power breaker cells create a border between two separate power domains, according to one embodiment.

[0019] Figure 8 is an illustration of an example arrangement, wherein power breaker cells provide discontinuities in metal straps, according to one embodiment.

[0020] Figure 9 is an illustration of a power breaker cell 901 implemented in a semiconductor device, according to one embodiment.

[0021] Figure 10 is an illustration of a flow diagram of an example method of designing a chip using power breaker cells, such as those of Figures 2-6, according to one embodiment.

[0022] Figure 11 is an illustration of a flow diagram of an example method of power collapsing sub cores, according to one embodiment.

DETAILED DESCRIPTION

[0023] Various embodiments provide for irregularly shaped sub cores within a core, where each of those sub cores is independently powered. An example is shown in Figure 1. This is in contrast to the strictly rectangular or rectilinear shapes of conventional sub cores.

[0024] The shapes of the sub cores are facilitated by the use of power breaker cells. The power breaker cells, during the design phase, include standard library cells that provide discontinuities in a power rail. For example, many power rails may traverse a dimension of the core and intersect multiple ones of the sub cores. At a boundary between one sub core and another sub core, the power breaker cells define that boundary at least at the power rails.

[0025] Each of the power breaker cells abuts two sub cores, and the sub cores are in separate power domains. The discontinuities provided by the power breaker cells allow each sub core to be powered separately. Accordingly, the processor core may power down one or more of the sub cores at a given time and independently of other sub cores.

[0026] Continuing with the example, various embodiments include designs for the power breaker cells themselves. One example embodiment includes a power breaker cell that is relatively narrow – either two or three grid units, though the scope of embodiments is not limited to any particular width. The power breaker cell includes two power rail contacts for VDD with a discontinuity between them, and it provides discontinuities within multiple metal layers. The power breaker cell also includes at

least one contact for ground or VSS. The power breaker cell may be implemented so that its two power rail contacts for VDD are built within the same N well. Furthermore, the power rail contacts may omit a physical electrical coupling to the N well within the boundaries of the power breaker cell.

[0027] Various embodiments may provide advantages over conventional solutions. For instance, a design having irregular boundaries for sub cores may reduce the space between those sub cores because the irregular shapes are not forced to fit within a prescribed rectangle or rectilinear shape. Thus, area of the chip can be used more efficiently.

[0028] Figure 1 is a simplified diagram illustrating an example processing core having multiple sub cores. The outside rectangle shape in Figure 1 represents the core, and the irregular shapes labeled sub core 0 … sub core 3 are the sub cores. Of note in Figure 1 is that the sub cores are not rectangular shaped, nor are they rectilinear shaped. Instead, the shapes of the sub cores simply conform to the boundaries of each other and to the core as a whole, thereby creating various irregular shapes. This is markedly different from conventional sub core layouts, wherein each of the sub cores is rectangular shaped or rectilinear shaped.

[0029] Continuing with the example, each of the different sub cores corresponds to a different power domain so that the sub cores are powered independently. In this example, the power domain boundaries conform to the irregular shapes of the sub core boundaries by use of standard cells that create a discontinuity in the power rails of the cores. For instance, power rail structures extend laterally (horizontal in this example) from one side of the core to the other and traversing multiple sub cores. Power rails are implemented in the metal layers of the semiconductor device, and a given rail may include a power and a ground or a power and a complementary power. Various

embodiments may place discontinuities in a power rail structure by including power breaker cells. An example of a power breaker cell is shown in Figure 2.

[0030] Figure 2 is a simplified diagram of a power breaker cell 200 adapted according to one embodiment. The power breaker cell 200 may be used to create discontinuities in the power rails at the boundaries of the sub cores of Figure 1.

[0031] Figure 2 includes a portion of a single power rail structure, and it is understood that the core of Figure 1 would include numerous power rail structures extending across its lateral dimension. The power breaker cell 200 is a standard cell, and it is at a level of granularity corresponding to a gate, such as an AND gate. When it is manufactured, power breaker cell 200 includes a deleted portion of metal 201 in the power rail structure. In the example of Figure 1, the irregular boundaries between one sub core and another sub core include power breaker cells that make power rails discontinuous at the boundaries. A given boundary may include hundreds or thousands of power breaker cells. In the example of Figure 2, the design would further include the left VDD rail 204 in communication with a first source of power and the right VDD rail 206 in communication with a second source of power so that each power domain may be turned on or off independently. As shown, the ground rail 208 is not broken by power breaker cell 200, however in some embodiments, other power breaker cells may break ground rails as well.

[0032] The design process for the core of Figure 1 is different from the conventional design process applied to conventional cores. For example, the design process for conventional cores would include rectilinear shapes being assigned to each of the different sub cores, with different teams working on each of the different sub cores. If an assigned shape prevents a particular sub core from complying with a constraint, then the rectilinear shape may be changed through an iterative process of

negotiation with teams for other sub cores and a core manager. By contrast, the design process for the core of Figure 1 dispenses with the assigned rectilinear shapes at the beginning and includes a team working to lay out the sub cores together within the large rectangular shape of the core itself. Power breaker cells may be placed late in the design process and may even be moved during later iterations of the design process as appropriate. This is described in more detail with respect to Figure 10. An advantage of the design process of the core of Figure 1 is that it may result in fewer iterations because the shapes of the individual sub cores are less constrained. Another advantage of the design process of the core of Figure 1 is that it may reduce the amount of empty space in boundaries between sub cores.

[0033] Figure 3 is an illustration of an example power breaker cell 300 adapted according to one embodiment and shown in more detail than is shown in Figure 2. Power breaker cell 300 is a 3-grid power breaker cell. VDD rail 310 is divided into two portions – a first VDD contact 310a on the left, and a second VDD contact 310b on the right, similar to rail portions 204 and 206 of Figure 2. Power breaker cell 300 includes a discontinuity 311 in the VDD rail 310, thereby separating power domains. The VSS rail 312 extends across the lateral dimension of the power breaker cell 300 without a discontinuity. The box labeled 320 is an N well (e.g., and N-type doped region of semiconductor on a P-type substrate), and both VDD contacts 310a and 310b are in the same N well 320. Furthermore, the power breaker cell 300 does not include well tying within the boundaries of the cell, as the VDD contacts 310a and 310b are not directly in communication with the doped semiconductor material of the N well.

[0034] Of course, the scope of embodiments is not limited to a 3-grid power breaker cell. For instance, Figure 4 is an illustration of an example 2-grid power breaker cell 400 according to one embodiment, and it is substantially similar to the power

breaker 300 cell of Figure 3, though including one fewer grid in its width. The 2-grid power breaker cell 400 also is implemented within a common N well and does not include well tying within the boundaries of the cell.

[0035] The power breaker cells 300 and 400 may be fabricated in any appropriate manner and using any appropriate materials. In one example, power breaker cells 300 and 400 include a floating gate 301. Floating gate 301 may include, for instance, a polysilicon gate that is not associated with any transistor. Diffusion breaks 302 are placed along a dimensional extent of the cell and may also be fabricated using polysilicon or other appropriate material to prevent diffusion of ions across the boundary of the power breaker cell. Items 303 include metal structures, which are similar to contacts that would be used at a source or a drain of a transistor, but in this example are not associated with any transistor. The VDD contacts 310a and 310b and VSS rail 312 are also metal, although they may be constructed in different layers than the metal structures 303.

[0036] As noted above, power breaker cells 200, 300, 400 are standard cells in a library. Accordingly, power breaker cells 200, 300, 400 may be embodied as data in a database, where the data describes how to create the cells in a semiconductor wafer. Of course, from another standpoint, power breaker cells 200, 300, 400 may be physically embodied on a semiconductor substrate, whether as part of a wafer or part of a diced chip. Figures 3 and 4 illustrate the physical features that may be used in various embodiments, whether described as part of a standard cell or physically implemented in silicon.

[0037] Power breaker cells 200, 300, 400 may be used within core designs to separate two sub cores in different power domains. Thus, a given power breaker cell may abut another standard cell on its left and yet another standard cell on its right,

where those other standard cells correspond to different sub cores. In such instances, VDD contact 310a would be in electrical communication with a VDD rail portion from the left, and VDD contact 310b would be in electrical communication with a VDD rail portion from the right. Similarly, VSS rail 312 would be in electrical contact with other VSS rail portions on both the left and the right. In this matter, a given power breaker cell provides an electrical discontinuity in a power rail at a boundary between two sub cores.

[0038] Figure 5 shows an example implementation, in which power breaker cell 400 abuts other cells 501 and 502. There is a discontinuity 510 in the VDD power rail 511, and VDD contacts 310a and 310b are electrically coupled with the metal of the VDD power rail 511. VSS power rail 512 is continuous for the entire width of the portion shown in Figure 5, including the portion that is attributable to VSS rail 312.

[0039] The cell 501 may be associated with a first sub core in a first power domain, and the cell 502 may be associated with a second sub core in a second power domain. Power breaker cell 400 abuts two power domains and provides the power rail discontinuity 510. Or put another way, power breaker cell 400 represents a border between two sub cores and two power domains. Both VDD power rail 511 and VSS power rail 512 span from one power domain to another and from one sub core to another, with power breaker cell 400 providing a boundary between the sub cores and power domains.

[0040] It is understood that a given sub core would include hundreds or thousands of other standard cells, and those additional standard cells are not shown herein. Nevertheless, the border between the two sub cores would extend both above and below the portion shown in Figure 5. And when viewed at a level of abstraction that allows viewing an entire sub core, the border between the two sub cores would include

hundreds or thousands of similar power breaker cells defining an irregular shape of the border.

[0041] Figure 6 illustrates an example implementation, wherein multiple power breaker cells create a border 601 between two separate power domains, domain A and domain B. An example of power breaker cell is power breaker cell 602, which may be similar to the power breaker cells 200, 300, 400 described above.

[0042] The horizontal lines in the Figure 6 represent VDD and VSS rails. For instance, VDD rail 611 is paired with VSS rail 612, and they are both associated with power breaker cell 602. Power breaker cell 602 provides a discontinuity in VDD rail 611 and may or may not also provide a discontinuity at VSS rail 612.

[0043] In this example, power is provided to the power domains A and B by metal straps 621-624 for distribution to the VDD rails. Figure 6 provides a top-down illustration, and it is understood that metal straps 621-624 are implemented in different metal layers than are the smaller metal lines exemplified by VDD rail 611 and VSS rail 612. The metal straps 621-624 tap to those progressively smaller metal lines exemplified by VDD rail 611 and VSS rail 612 using vias (not shown). Metal straps 621-624 are interconnected to provide a low resistance path for current and to allow for redundancy. For instance, metal straps 621 and 622 belong to power domain A, and they are interconnected by virtue of vias (not shown) that reach different metal layers and couple to smaller metal lines exemplified by VDD rail 611 and VSS rail 612. Similarly, metal straps 624 and 623 belong to power domain B, which are also interconnected by virtue of vias that coupled to the smaller metal lines exemplified by VDD rail 611 and VSS rail 612.

[0044] The power breaker cells provide discontinuities at their respective VDD rails (and possibly VSS rails) in one or more metal layers. For instance, in designs that

distribute power using two or more metal layers, the power breaker cells may provide discontinuities in those two or more metal layers.

[0045] The border 601 between domains A and B takes an irregular shape at this level of abstraction. When looking at the scale of gates themselves, irregularly shaped may appear to be rectilinear, but looking at a scale that is an order of magnitude larger or two orders of magnitude larger than a given gate, the border 601 appears to be irregularly shaped. The view of the boundaries between the sub cores in Figure 1 highlights the irregular shape. In fact, the view of the border 601 in Figure 6 may be only a small section of one of the exemplary boundaries of Figure 1.

[0046] Although not shown in Figure 6, power domain A and power domain B may be in communication with separate power supplies in this embodiment. For instance, a power supply may include a Power Management Integrated Circuit (PMIC), and in some embodiments metal straps 621 and 622 may be in communication with a first PMIC, whereas metal straps 623 and 624 may be in communication with a second PMIC. This is one example of how two different power domains may coexist within a semiconductor device. In another example, a single power supply may provide power to both power domains A and B, and a power distribution network including switches may be implemented between the power supply and the straps 621-624. The switches may be opened and closed to separately power each power domain. The scope of embodiments is not limited to any particular power scheme for power distribution network.

[0047] Various embodiments include power collapsing power domains independently of each other. For instance, some power saving techniques include reducing a voltage or removing power altogether to a given sub core (power collapsing) when that sub core is not in use. For instance, when a particular camera sub cores not in use, a power supply may discontinue providing power to that particular sub core, while

at the same time providing power to other sub cores, assuming those sub cores are in different power domains. The power breaker cells of Figures 2-6 allow for separation of the power domains, thereby allowing sub cores to be powered independently of each other, including power collapsing sub cores independently of each other.

[0048] Figure 7 provides an illustration of another border 701 between power domain A and power domain B, according to one embodiment. For instance, the section of the border 701 shown in Figure 7 may be part of a larger border that includes the portion shown in Figure 6 as well, perhaps vertically above or below the view of Figure 7.

[0049] Furthermore, other embodiments may include additionally or alternatively using power breaker cells in metal straps, such as metal straps 621-624. Figure 8 is an illustration of an example arrangement, wherein power breaker cells 811 and 812 provide discontinuities in metal straps 821 and 822. Taken together, the embodiments of Figures 6-8 illustrate that the power breaker cells may be used at any appropriate structure to create a discontinuity in a power or ground rail.

[0050] Figure 9 is an illustration of a power breaker cell 901 implemented in a semiconductor device, according to one embodiment. Figure 9 is a cutaway view from a side of the semiconductor device, showing two metal layers, Metal Y and Metal X disposed within layers of dielectric. Power breaker cell 901 has floating gate 910, and it includes portions of metal that are in Metal Y and Metal X layers. Figure 9 illustrates that a power breaker cell, as implemented in a semiconductor device, provides discontinuities in one or more metal layers of the semiconductor device. When implemented as a standard cell in a library, the power breaker cell 901 may mark various metal layers as XOR to indicate that those metal layers have discontinuities within the cell. While Figure 9 shows a discontinuity in two metal layers, it is

understood that various semiconductor devices may have a multitude of metal layers, and a given power breaker cell may provide discontinuities in one or more of the metal layers as appropriate to separate power domains.

[0051] Figure 10 is a flow diagram of an example method 1000 according to one embodiment. Method 1000 includes a design flow, which may be used by a design team, including a design team for an individual core of the chip as well as a design team for the chip as a whole. Much of the actions described with respect to Figure 10 may be performed using Electronic Design Automation (EDA) software. The EDA software treats the standard cells as logical entities that describe a physical part of a semiconductor device. During the design process, a team member may select standard cells from a library and instruct the EDA software to place those standard cells at specific places within the design of the semiconductor device.

[0052] The product of the EDA software includes one or more computer readable files that describe the structure of the semiconductor device. The one or more computer readable files may be read by tools at a foundry and used during the fabrication process to produce physical semiconductor devices based on the design.

[0053] At action 1010, there is initial placement optimization for the individual sub cores. For instance, one or more members of a design team may start with a rough idea about initial placement of the sub cores within the core. For instance, in the design of Figure 1, sub core zero is at the top right portion of the core, and the other sub cores have their respective places. Such initial placements may be based on prior knowledge of the team regarding general physical shapes of the logic and general timing constraints.

[0054] At action 1020, the team designs the sub cores according to their initial placements. Of note at action 1020, the individual sub cores are not assigned to

rectilinear shapes. Rather they are merged into a single plan and implemented at a same level. Action 1020 may include iterations to ensure that logic is placed appropriately close to its parent logic and that other design constraints are met. Furthermore, power domain separation is respected, so that portions of logic that belong to a same power domain are placed contiguous with other logic of that same power domain.

[0055] Conventional design processes, which would assign sub cores to respective rectilinear shapes, would typically use a greater number of iterations due to fitting the various sub cores to their rectilinear shapes. However, action 1020 eliminates iterations that would otherwise be attributable to fitting the various sub cores into respective rectilinear shapes. The initial placement and the design of the sub cores may be performed by a team or by an individual with assistance from the EDA software. In this example action 1020 omits the conventional practice of dividing each sub core to a separate team.

[0056] Once the sub cores are designed and have separate power domain boundaries, action 1030 introduces power breaker cells at the boundaries of the power domains, which are expected to be the same as the boundaries of the sub cores. Of note at action 1030, the power breaker cells are introduced toward the end of the design process for the core. An advantage of various embodiments is that the power breaker cells are standard cells that can be placed easily and conveniently at this stage of the design process and would not be expected to result in substantial redesign. Continuing with action 1030, any further incremental placement optimization is restricted to within the different power domains, although further iterations requiring changes to the boundaries may be achieved by moving the power breaker cells that Mark the boundaries.

[0057] At action 1040, it is assumed that the design for the core is substantially complete and should be assembled together with designs for other cores of the chip as a whole. Accordingly, the data of the core design is submitted to the chip team. The chip team ensures that chip-level constraints are met, such as power and timing constraints. It is possible that action 1040 may result in one or more iterations.

[0058] Action 1050 includes tape out. Tape out may involve sending one or more design files to a foundry.

[0059] The scope of embodiments is not limited to the specific series of actions in Figure 10. Rather, various embodiments may add, omit, rearrange, or modify one or more actions. For instance, more or fewer iterations may be used during steps 1020-1040 as appropriate.

[0060] Various embodiments may include one or more advantages over conventional processes. For instance, to the extent that method 1000 achieves a fewer number of iterations than a conventional design process, that reduction in man-hours means more efficient design and perhaps cost savings. Additionally, boundaries between the different power domains may correspond to the width of a single power breaker cell. This is in contrast to conventional design processes that assign individual sub cores to respective rectilinear shapes, where those rectilinear shapes may be separated by unused space on the chip. Various embodiments herein eliminate wasteful space between sub core rectilinear shapes, thereby reducing semiconductor area. Additionally, embodiments described herein preserve the ability to separately power collapse different sub cores while at the same time minimizing wasted semiconductor area.

[0061] A flow diagram of an example method 1100 for separately power collapsing sub cores is illustrated in Figure 11. In one example, method 1100 is

performed by a power management logic in a core. For instance, a camera core or other processor core may include power management logic that shuts off power to sub cores during times when those sub cores are not in use. The logic may execute computer readable instructions to perform power collapsing. Power collapsing the sub cores may reduce power consumption that is due to leakage current of individual transistors. Furthermore, power collapsing may be performed as an alternative to or in addition to other power management techniques, such as clock gating and the like.

[0062] Method 1100 of Figure 11 is performed in a core, such as that shown in Figure 1 having multiple sub cores and irregular boundaries between power domains. The boundaries are provided by power breaker cells, such as those illustrated in Figures 2-9. Furthermore, while method 1000 of Figure 10 (described above) addresses power breaker cells as standard cells used during a design process, method 1100 of Figure 11 is performed with respect to a chip that has been fabricated and is in use either during testing or during normal operation of the device.

[0063] At action 1110, the power management logic power collapses a first sub core while a second sub core is powered. In other words, the power management logic may temporarily stop powering the first one of the sub cores while continuing to provide power to the second one of the sub cores. The two sub cores in this example are adjacent each other and are separated along their boundary by power breaker cells. Despite the discontinuities provided by the power breaker cells, multiple power rails traverse a dimension the core from the first sub core to the second sub core. A given power rail having a discontinuity may provide an operating voltage to the second core while not being powered through its portion within the first core.

[0064] At action 1120, the power management logic power collapses the second sub core. At this point, the power management logic may provide power to the first sub

core or the first sub core may continue to be power collapsed. In any event, the first and second sub cores are individually power collapsible.

[0065] The scope of embodiments is not limited to the specific method shown in Figure 11. Other embodiments may add, omit, rearrange, or modify one or more actions. For instance, method 1100 may be performed multiple times as the core operates, so that each one of the sub cores may be power collapsed or maybe powered as appropriate. Both of the sub cores may be on at the same time, both of the sub cores may be power collapsed at the same time, or one of the sub cores may be power collapsed as the other receives power. Additionally, method 1100 is described with respect to two sub cores, but it is understood that various embodiments may include two or more sub cores that are individually power collapsible. The scope of embodiments is not limited to any particular number of sub cores. Moreover, the power collapsing and powering up of sub cores of method 1100 may be performed using any appropriate technique, such as opening or closing switches between power supplies and the individual sub cores, turning power supplies on or off, and the like.

[0066] As those of some skill in this art will by now appreciate and depending on the particular application at hand, many modifications, substitutions and variations can be made in and to the materials, apparatus, configurations and methods of use of the devices of the present disclosure without departing from the spirit and scope thereof. In light of this, the scope of the present disclosure should not be limited to that of the particular embodiments illustrated and described herein, as they are merely by way of some examples thereof, but rather, should be fully commensurate with that of the claims appended hereafter and their functional equivalents.

CLAIMS

What is claimed is:

1. A semiconductor device comprising:

a processing core having a plurality of sub cores;

a plurality of power rails spanning from a first sub core to a second sub core of the plurality of sub cores, the plurality of power rails configured to provide an operating voltage to each of the first sub core and the second sub core; and

a plurality of cells defining a boundary between the first sub core and the second sub core, each of the cells providing a discontinuity in a respective power rail, wherein the discontinuity includes a break in the respective power rail in more than one layer of the semiconductor device.

2. The semiconductor device of claim 1, wherein each of the cells comprises:

a first VDD contact; and

a second VDD contact, the first and second VDD contacts being built within a common N-type doped region of a substrate of the semiconductor device.

3. The semiconductor device of claim 2, wherein the first VDD contact is coupled to the respective power rail on a side of the boundary corresponding to the first sub core, and wherein the second VDD contact is coupled to the respective power rail on a side of the boundary corresponding to the second sub core.

4. The semiconductor device of claim 2, wherein the first VDD contact and the second VDD contact are not directly coupled with doped semiconductor material of the N-type doped region.

5. The semiconductor device of claim 1, wherein each power rail includes a conductive line for VSS and a conductive line for VDD, wherein the conductive line for VDD include the discontinuity.

6. The semiconductor device of claim 1, wherein each of the cells comprises:

a floating gate structure; and

a plurality of diffusion breaks along an extent of the cell.

7. The semiconductor device of claim 1, further comprising:

an additional power rail perpendicular to the plurality of power rails and coupled with the plurality of power rails within the first sub core; and

an additional cell providing a discontinuity in the additional power rail.

8. The semiconductor device of claim 1, wherein each of the cells has a width of three grid units or less.

9. The semiconductor device of claim 1, wherein the first sub core and the second sub core are in separate power domains, and wherein the separate power domains correspond to separate power supplies.

10. A semiconductor chip comprising:

a plurality of processing cores disposed within semiconductor material of the semiconductor chip, a first one of the cores having a first sub core and a second sub core;

a power rail spanning from the first sub core to the second sub core, the power rail including conductive lines within the semiconductor material; and

a cell abutting the first sub core and the second sub core on the power rail, the

cell providing a discontinuity in the power rail at a plurality of metal layers of the conductive lines.

11. The semiconductor chip of claim 10, wherein the cell comprises:

a first VDD contact; and

a second VDD contact, the first and second VDD contacts being built within a common N-type doped region of a substrate of the semiconductor chip.

12. The semiconductor chip of claim 11, wherein the first VDD contact is coupled to the power rail at the first sub core, and wherein the second VDD contact is coupled to the power rail at the second sub core.

13. The semiconductor chip of claim 11, wherein the first VDD contact and the second VDD contact are not directly coupled with doped semiconductor material of the N-type doped region.

14. The semiconductor chip of claim 10, wherein the power rail includes a conductive line for VSS and a conductive line for VDD, wherein the conductive line for VSS runs continuously through the cell.

15. The semiconductor chip of claim 10, wherein the first sub core and the second sub core are in separate power domains.

16. The semiconductor chip of claim 10, further comprising:

an additional power rail perpendicular to the power rail and in communication with the power rail within the first sub core; and

an additional cell providing a discontinuity in the additional power rail.

17. The semiconductor chip of claim 10, wherein the cell comprises:

a gate structure disposed on a substrate of the semiconductor chip and not associated with a transistor.

18. A semiconductor chip having a processing core, the processing core comprising:

a first sub core;

a second sub core;

means for distributing power to the first sub core and the second sub core; and

means for providing discontinuities in the power distributing means at multiple metal layers of the semiconductor chip, wherein a boundary between the first sub core and the second sub core is defined by the discontinuity providing means abutting the first core and the second core.

19. The semiconductor chip of claim 18, wherein the discontinuity providing means includes a plurality of cells at the boundary, wherein each of the cells comprises:

a first VDD contact; and

a second VDD contact, the first and second VDD contacts being built within a common N-type doped region of a substrate of the semiconductor chip.

20. The semiconductor chip of claim 19, wherein the first VDD contact is coupled to a respective power rail on a side of the boundary corresponding to the first sub core, and wherein the second VDD contact is coupled to the respective power rail on a side of the boundary corresponding to the second sub core.

21. The semiconductor chip of claim 19, wherein the first VDD contact and the second VDD contact are not directly coupled with doped semiconductor material of the N-type doped region.

22. The semiconductor chip of claim 18, wherein the power distributing means comprise a plurality of power rails, and wherein the discontinuity providing means comprise a plurality of cells defining the boundary at points where the power rails span from the first sub core to the second sub core.

23. The semiconductor chip of claim 22, wherein each power rail includes a conductive line for VSS and a conductive line for VDD, wherein the conductive lines for VSS do not include discontinuities provided by the cells.

24. The semiconductor chip of claim 18, wherein the processing core comprises a Graphics Processing Unit (GPU), and wherein the first sub core comprises a shader unit, and wherein the second sub core comprises an accumulator unit.

25. The semiconductor chip of claim 18, wherein the first sub core and the second sub core are in separate power domains, and further wherein the separate power domains correspond to separate power supplies.

26. A method comprising:

power collapsing a first sub core while providing power to a second sub core, wherein the first sub core and the second sub core are implemented in a processing core on a semiconductor chip, wherein a boundary between the first sub core and the second sub core includes a plurality of cells providing discontinuities in a plurality of power rails at multiple layers of the semiconductor chip; and

power collapsing the second sub core independently of power collapsing the first sub core.

27. The method of claim 26, further comprising power collapsing the first sub core and the second sub core at same time.

28. The method of claim 26, further comprising providing power to the second sub core and the first sub core at a same time.

29. The method of claim 26, wherein each of the cells comprises:

a first VDD contact; and

a second VDD contact, the first and second VDD contacts being built within a common N-type doped region of a substrate of the semiconductor device.

30. The method of claim 26, wherein the first VDD contact and the second VDD contact are not directly coupled with doped semiconductor material of the N-type doped region.

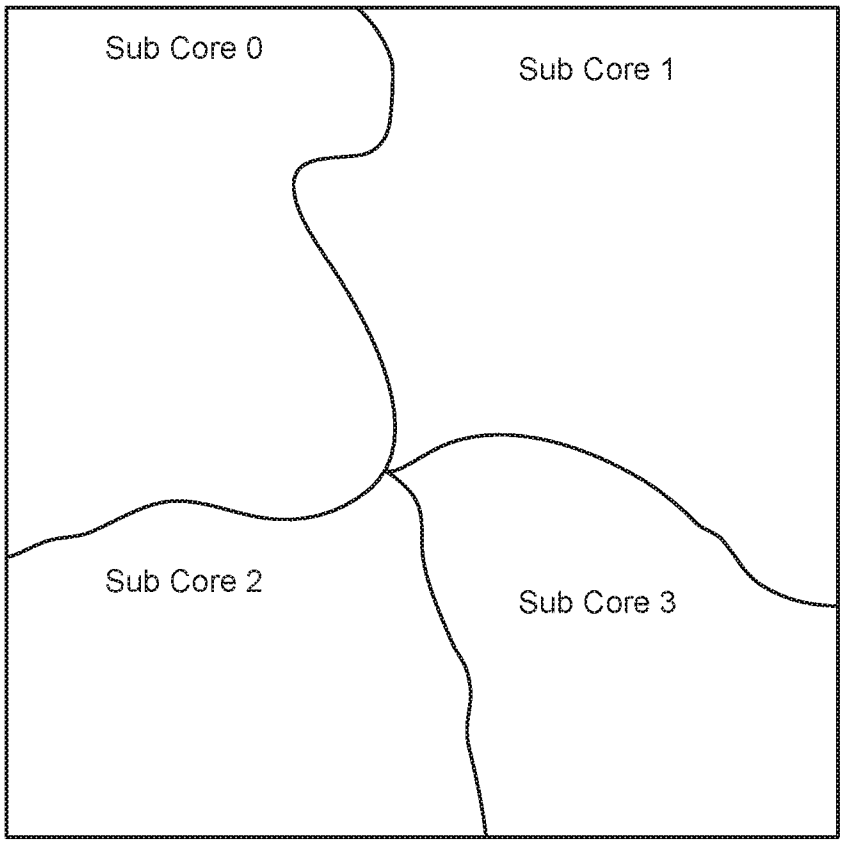

FIG. 1

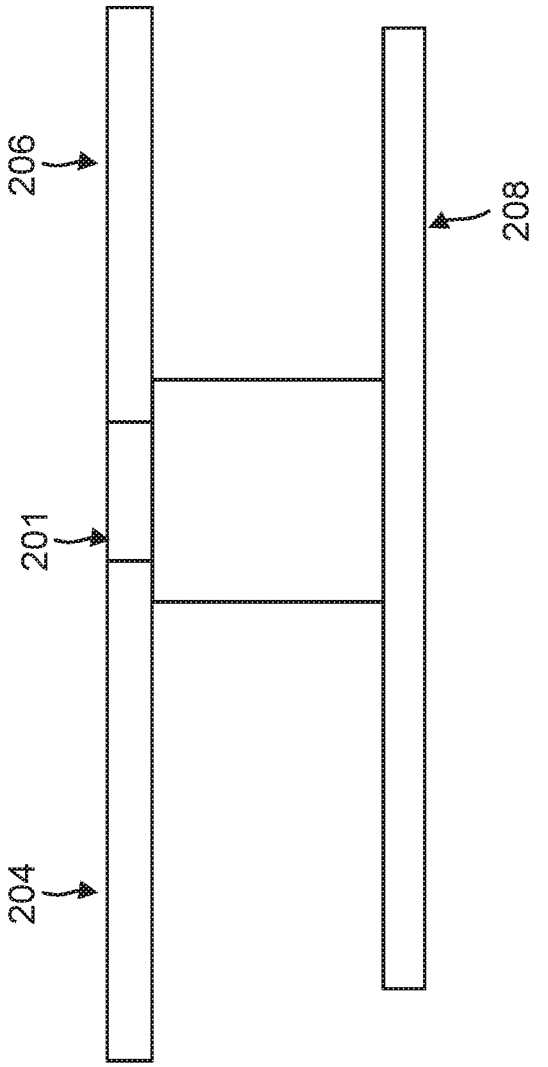

FIG. 2

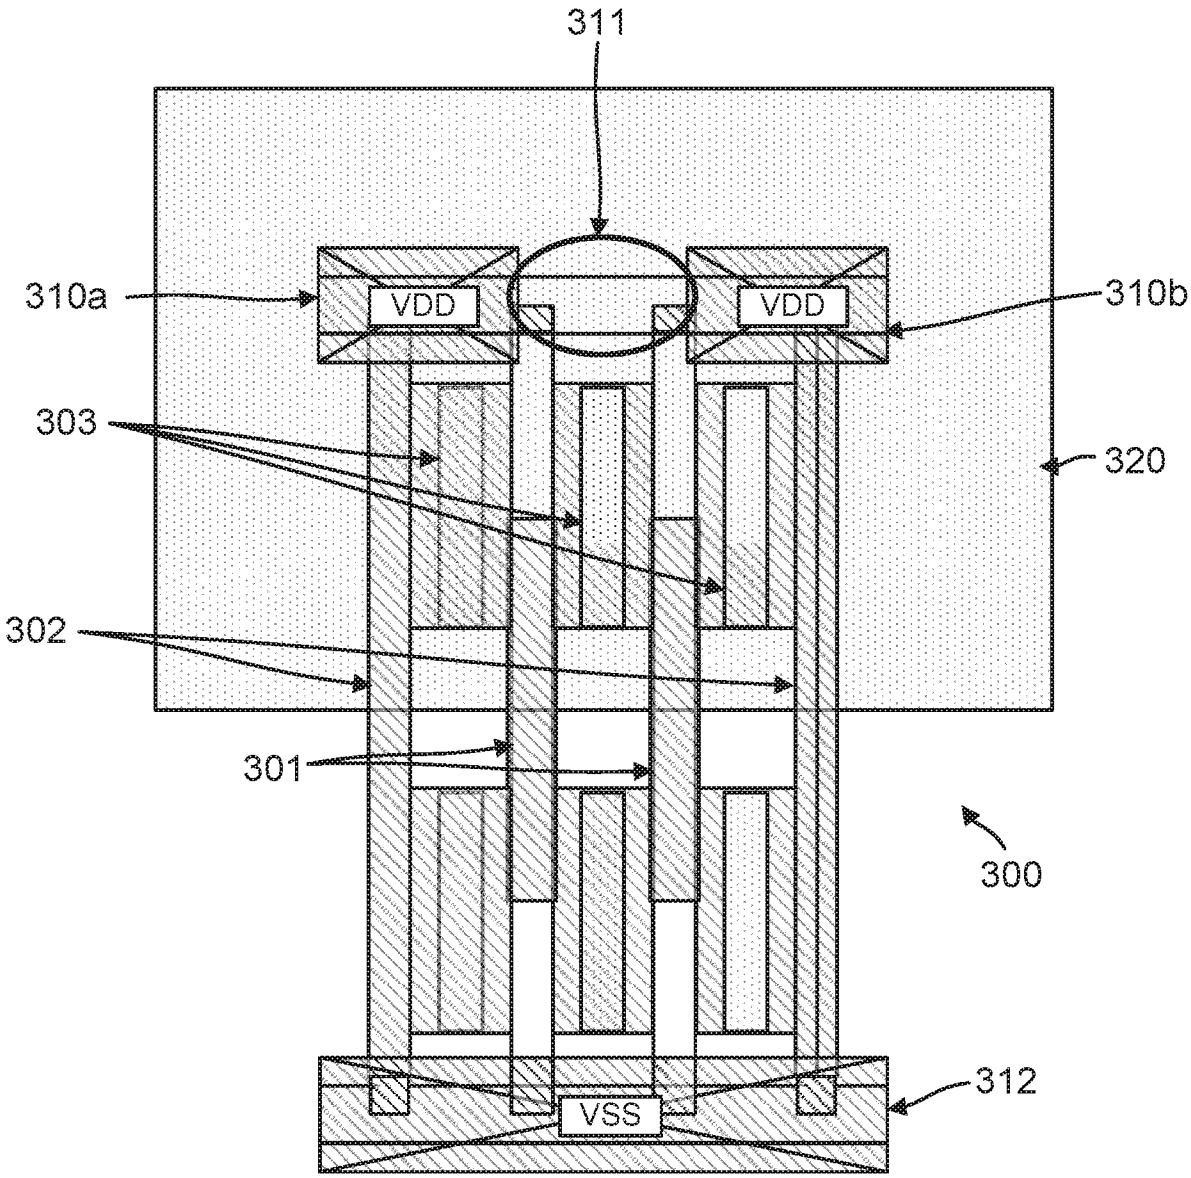

FIG. 3

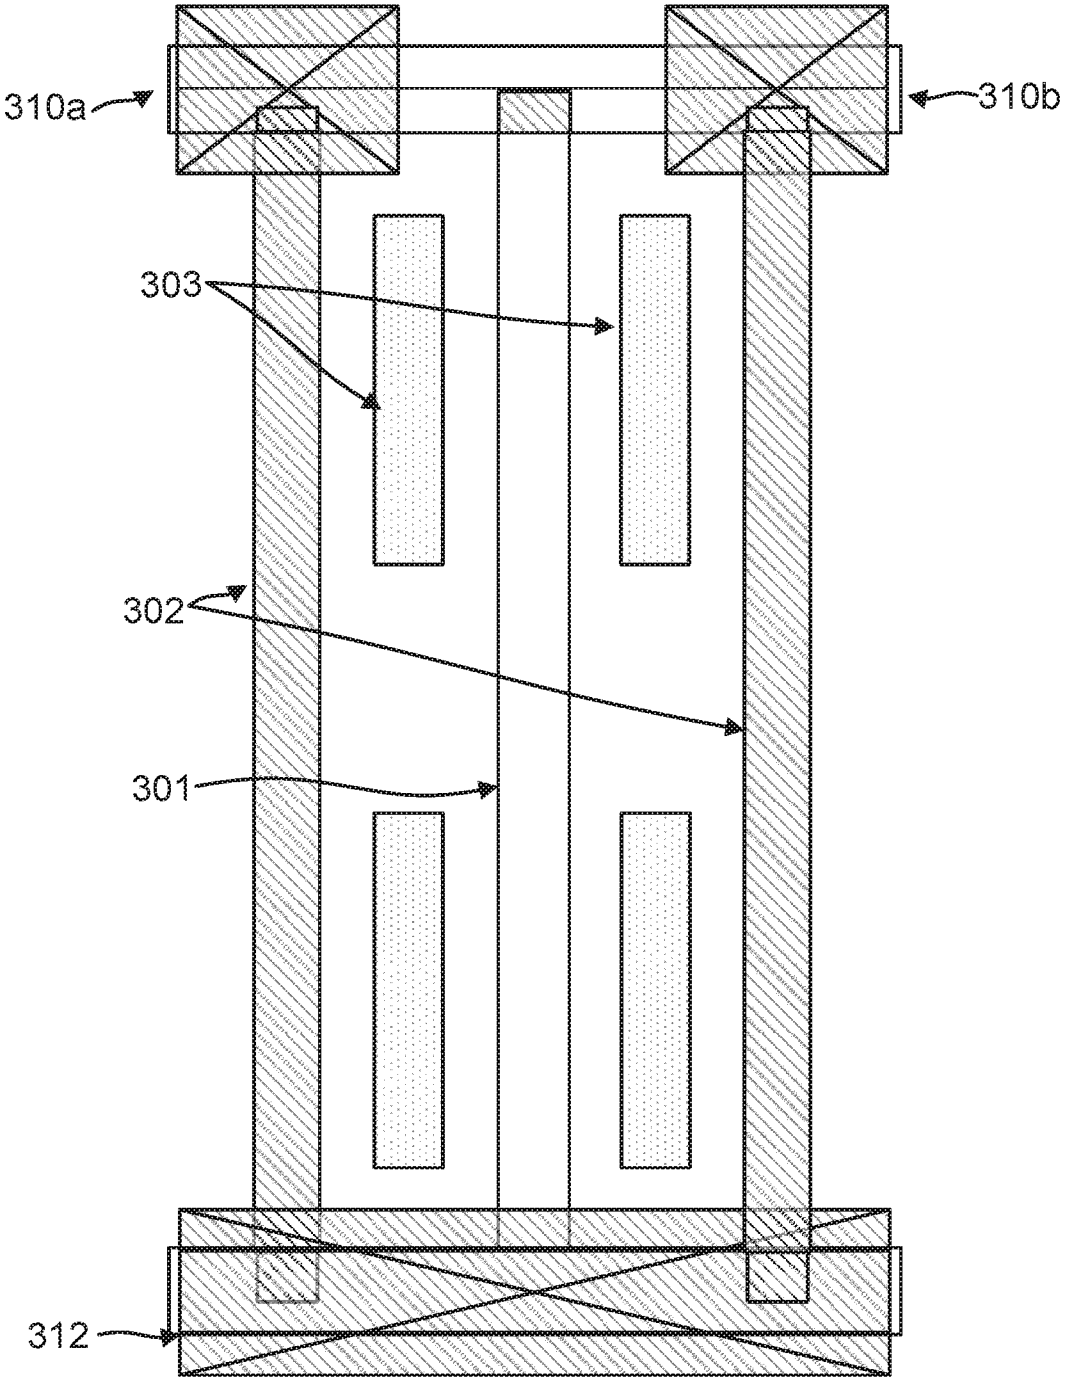

FIG. 4

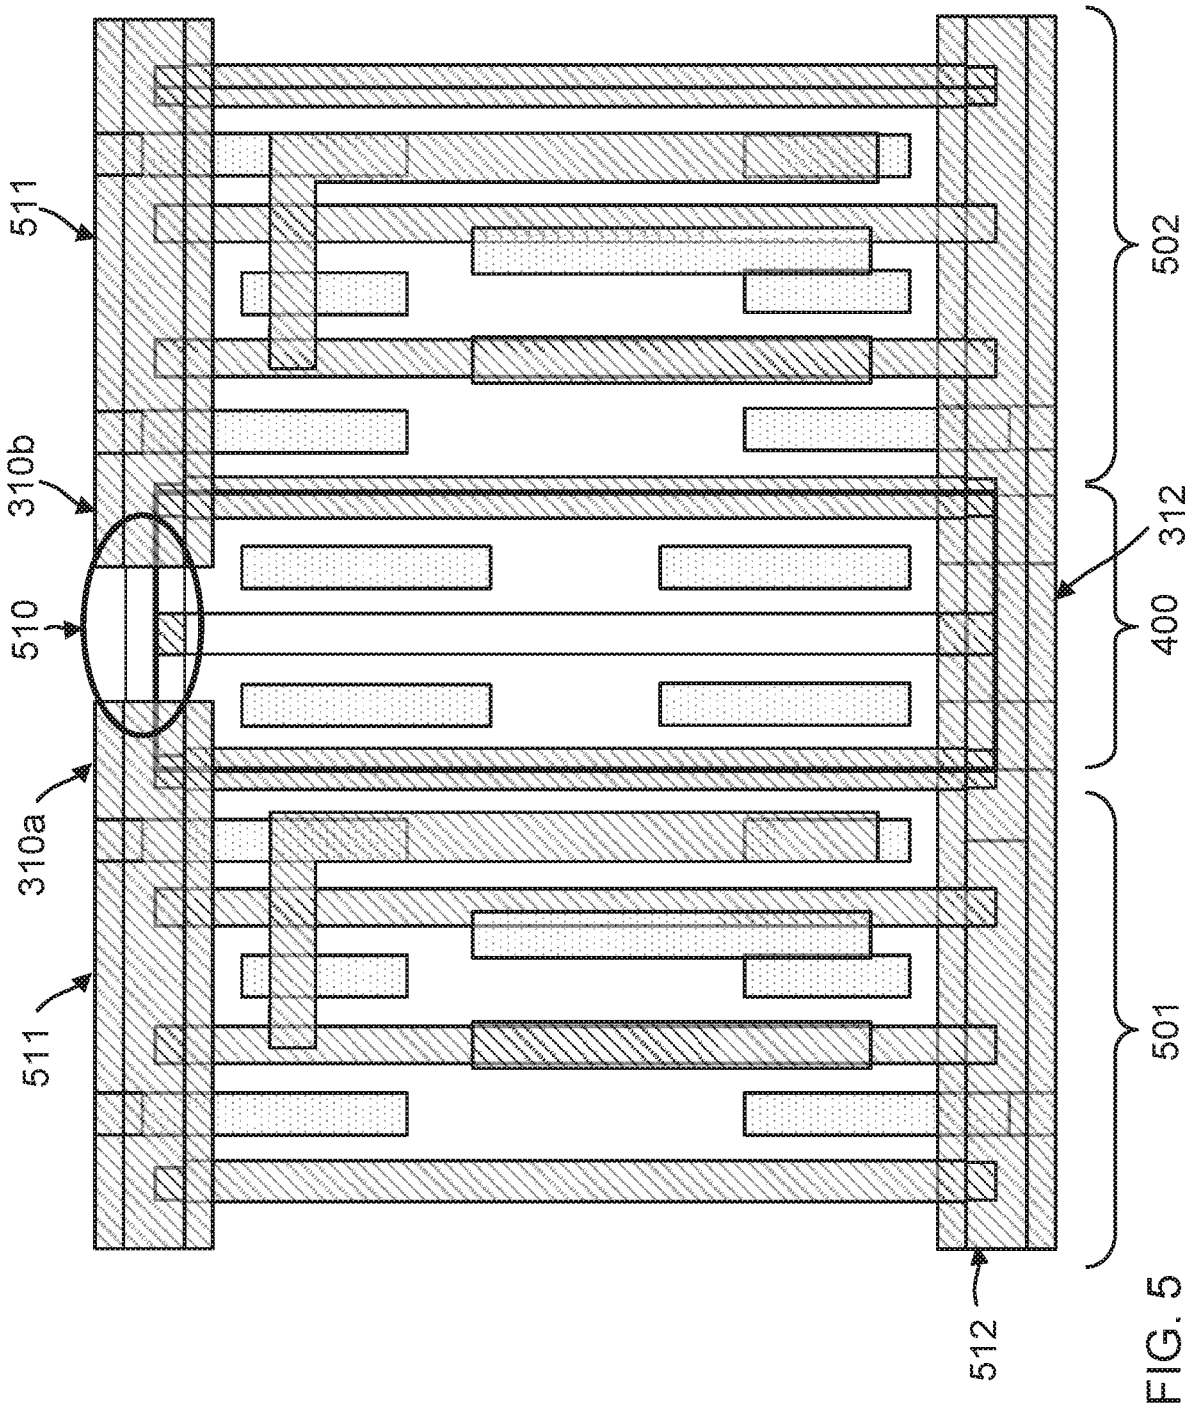

FIG. 5

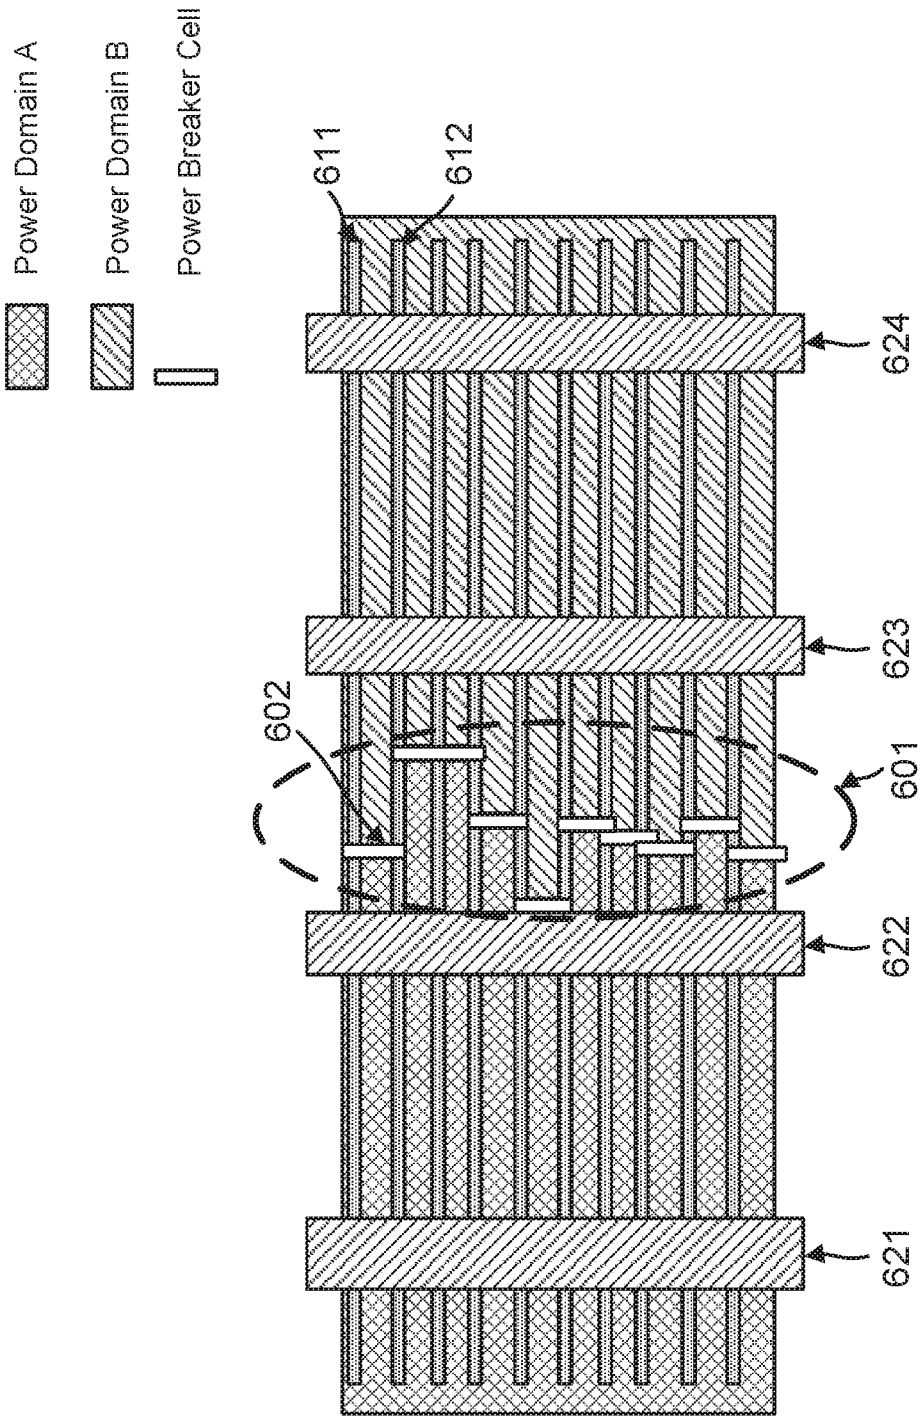
Power Domain A
Power Domain B
Power Breaker Cell
611
612
624
623
602
601
622
621

FIG. 6

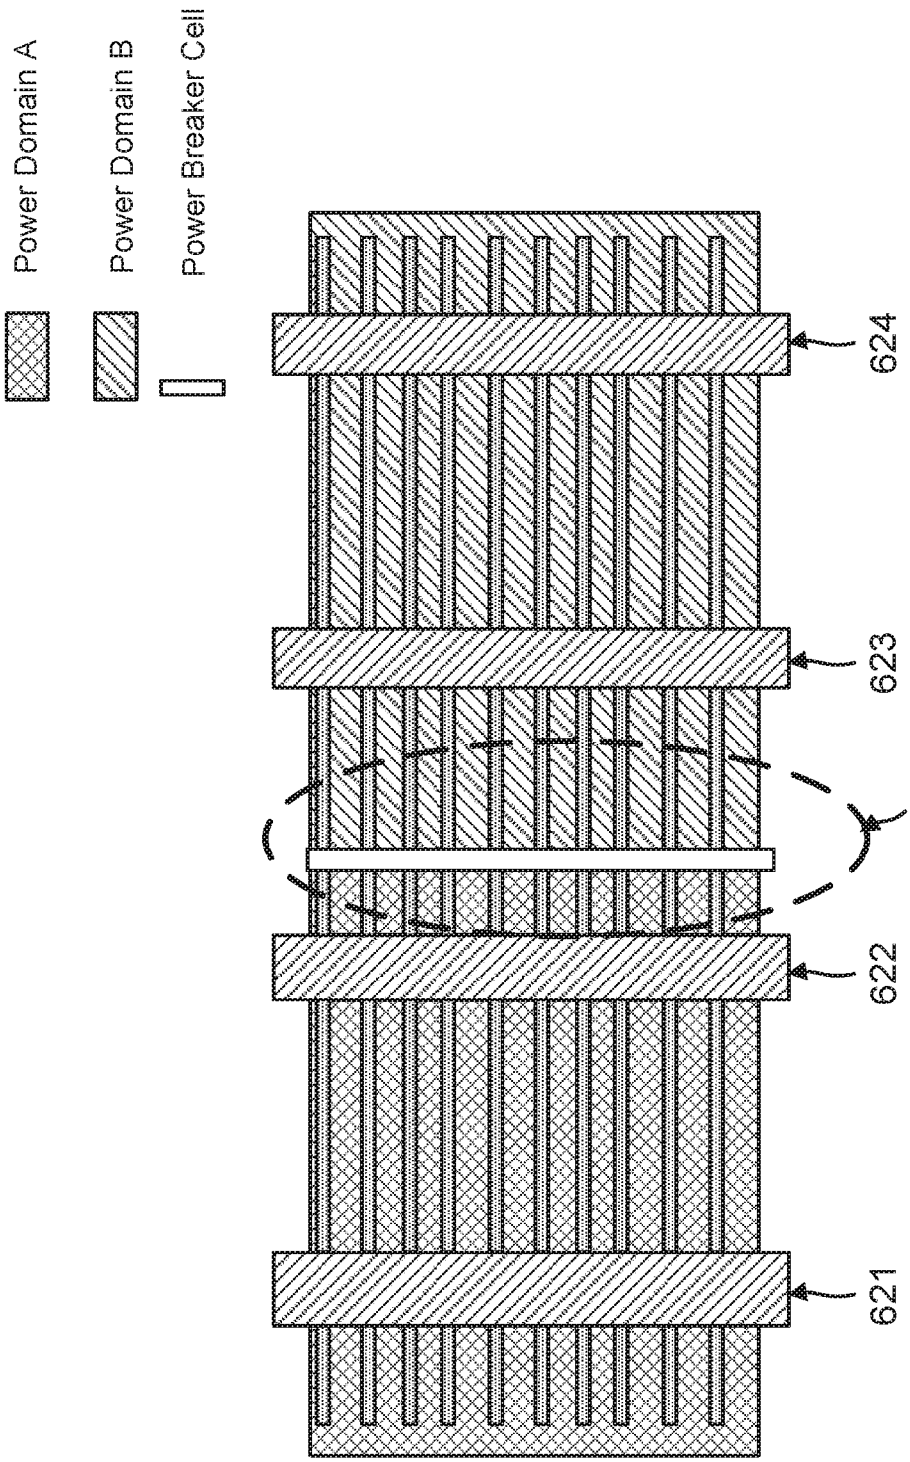
Power Domain A
Power Domain B
Power Breaker Cell
624
623
622
621
701

FIG. 7

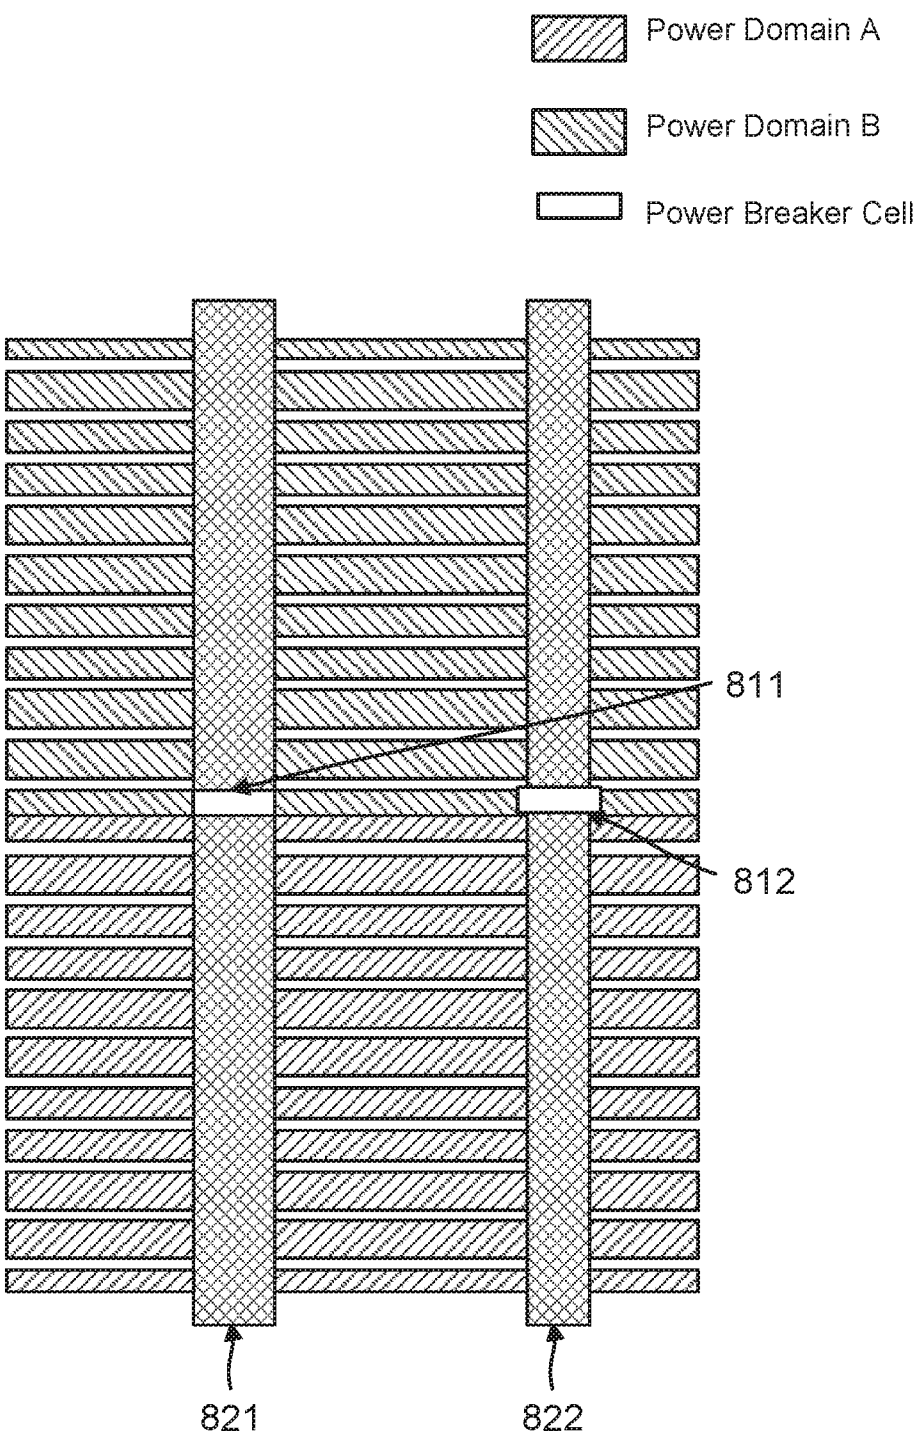
Power Domain A
Power Domain B
Power Breaker Cell
811
812
821
822

FIG. 8

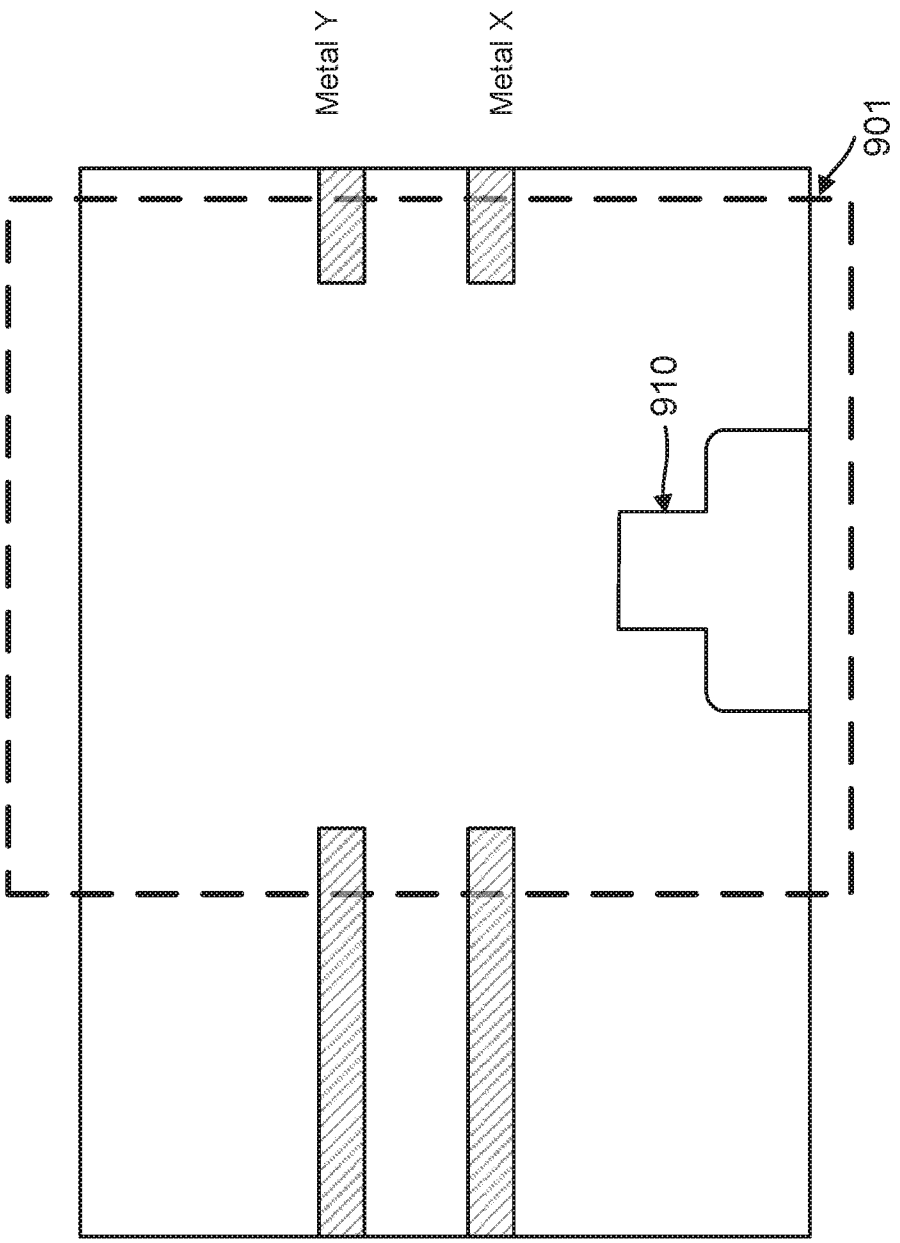

FIG. 9

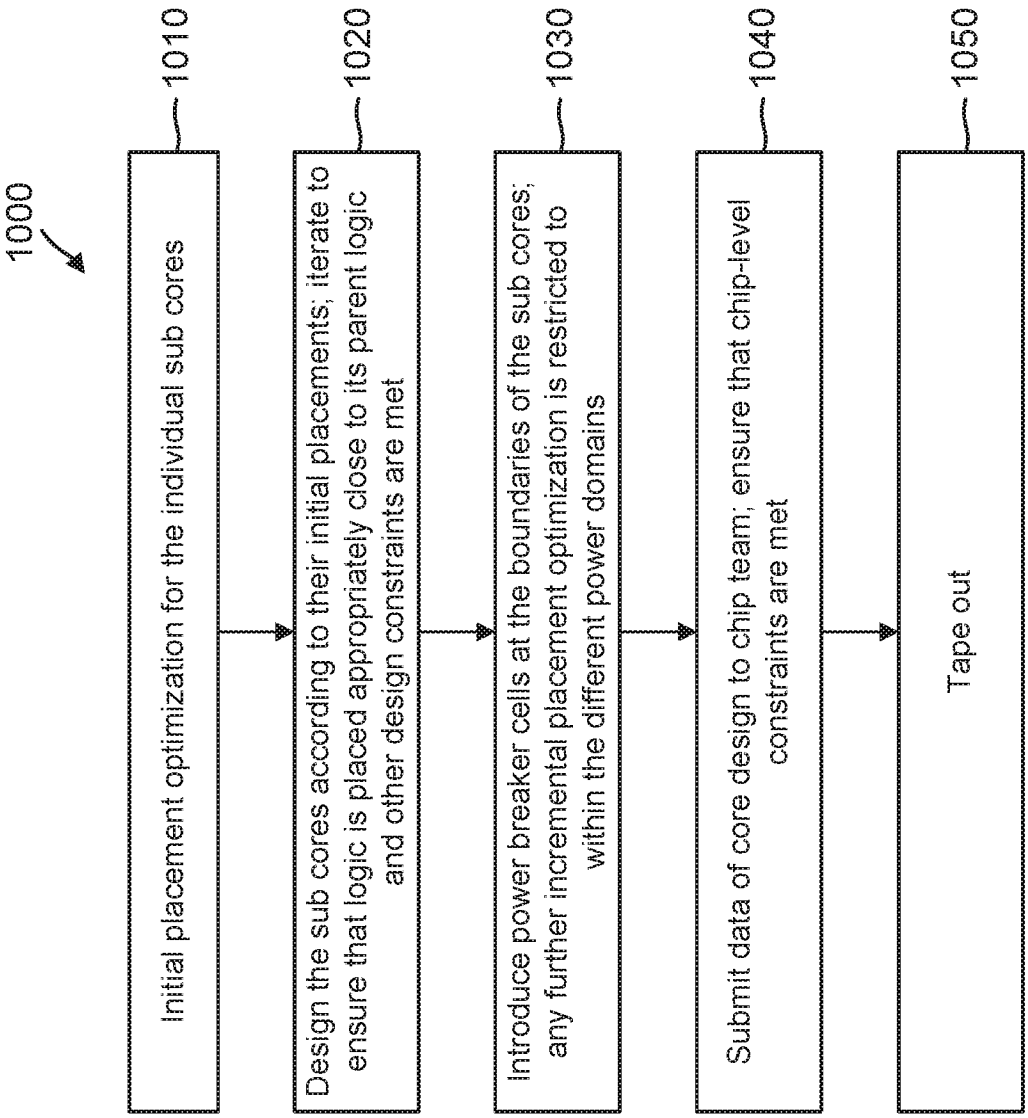
1000
1010
Initial placement optimization for the individual sub cores
1020
Design the sub cores according to their initial placements; iterate to ensure that logic is placed appropriately close to its parent logic and other design constraints are met
1030
Introduce power breaker cells at the boundaries of the sub cores; any further incremental placement optimization is restricted to within the different power domains
1040
Submit data of core design to chip team; ensure that chip-level constraints are met
1050
Tape out

FIG. 10

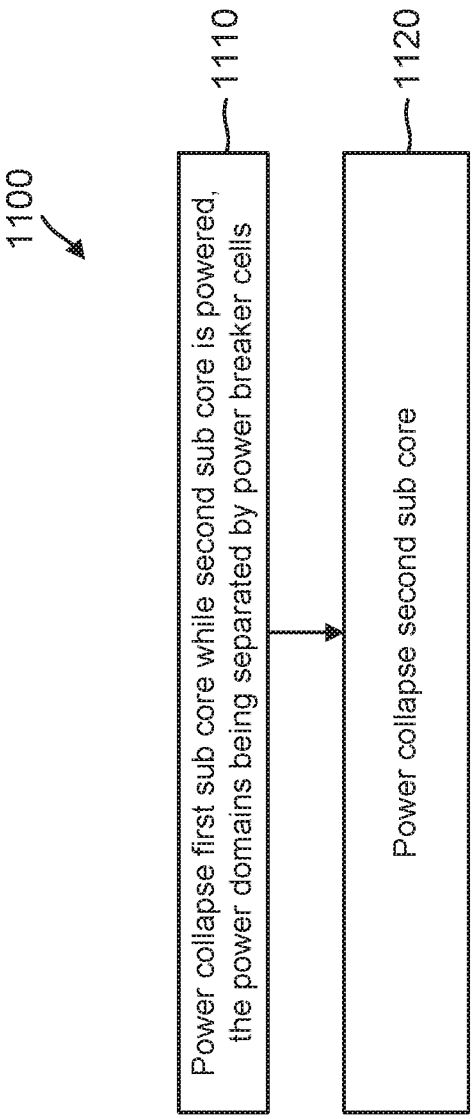
1100
1110
Power collapse first sub core while second sub core is powered, the power domains being separated by power breaker cells
1120
Power collapse second sub core

FIG. 11

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/027968

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F1/32
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F H01L G11C

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Anonymous: "Encounter User Guide, Product Version 7.1.2. Chapter 10: Low Power Design", , 30 May 2008 (2008-05-30), pages 367-432, XP055395027, Retrieved from the Internet: URL:http://w1npp.org/wx1gyx.org/topics/SKYWAR~1/LDSPED/565893~2.PDF [retrieved on 2017-08-01] page 368 - page 369 page 375 - page 376 page 383 ----- -/--	1-9, 26-30

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
2 August 2017	09/08/2017

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Alonso Nogueiro, L

2

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/027968

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☒ Claims Nos.: 10-25
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
see FURTHER INFORMATION sheet PCT/ISA/210

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.

☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.

☐ No protest accompanied the payment of additional search fees.

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2013/015882 A1 (DATTA ANIMESH [US] ET AL) 17 January 2013 (2013-01-17) paragraph [0031] - paragraph [0038] paragraph [0041] - paragraph [0045] figures 1-4	6
A	US 2004/036363 A1 (SARIG EREZ [IL] ET AL) 26 February 2004 (2004-02-26) paragraph [0029] - paragraph [0032] figure 5	4,30
A	US 2001/036108 A1 (TAKANO YOSHINORI [JP]) 1 November 2001 (2001-11-01) paragraph [0004] - paragraph [0008] paragraph [0109] - paragraph [0111] figures 13-15	6
A	US 4 627 152 A (DEHOND MITCHELL R [US] ET AL) 9 December 1986 (1986-12-09) column 1, line 40 - column 2, line 20 figure 2	6

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2017/027968

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013015882 A1	17-01-2013	CN 103765581 A	30-04-2014
		EP 2732467 A1	21-05-2014
		JP 5940660 B2	29-06-2016
		JP 2014527287 A	09-10-2014
		JP 2016171335 A	23-09-2016
		KR 20140046023 A	17-04-2014
		US 2013015882 A1	17-01-2013
		WO 2013010043 A1	17-01-2013
US 2004036363 A1	26-02-2004	NONE	
US 2001036108 A1	01-11-2001	JP 2001312893 A	09-11-2001
		US 2001036108 A1	01-11-2001
US 4627152 A	09-12-1986	CA 1241115 A	23-08-1988
		DE 3688146 D1	06-05-1993
		DE 3688146 T2	07-10-1993
		EP 0202535 A2	26-11-1986
		JP H07107918 B2	15-11-1995
		JP S61272948 A	03-12-1986
		US 4627152 A	09-12-1986

Continuation of Box II.2

Claims Nos.: 10-25

The present application contains 30 claims, of which 3 are independent claims in the system category, i.e. claims 1, 10 and 18. There is no clear distinction between independent claims 1, 10 and 18 because of overlapping scope and these claims are drafted in such a way that the set of claims as a whole is not in compliance with the provisions of clarity and conciseness of Article 6 PCT, as it is particularly burdensome for a skilled person to establish the subject-matter for which protection is sought. The non-compliance with the substantive provisions is to such an extent, that the search was performed taking into consideration the non-compliance in determining the extent of the search (PCT Guidelines 9.19 and 9.25).
The search was based on the subject-matter that, as far as can be understood, could reasonably be expected to be claimed later in the procedure, and the corresponding claims 1 to 9 and 26 to 30.

The applicant's attention is drawn to the fact that claims relating to inventions in respect of which no international search report has been established need not be the subject of an international preliminary examination (Rule 66.1(e) PCT). The applicant is advised that the EPO policy when acting as an International Preliminary Examining Authority is normally not to carry out a preliminary examination on matter which has not been searched. This is the case irrespective of whether or not the claims are amended following receipt of the search report or during any Chapter II procedure. If the application proceeds into the regional phase before the EPO, the applicant is reminded that a search may be carried out during examination before the EPO (see EPO Guidelines C-IV, 7.2), should the problems which led to the Article 17(2) declaration be overcome.