

TITLE: A semi-conductor integrated circuit

DESCRIPTION

The present invention relates to a semi-conductor integrated circuit, and especially, but not
5 exclusively, to a semi-conductor integrated circuit of the type comprising configurable logic circuit arrays.

A semi-conductor integrated circuit comprises a plurality of resources which can be divided up into routing resources, function resources and
10 storage/program resources. In our British Patent Specification No. 9410980.8 (which is itself a development of our British Patent No. 2180382 and U.S. Patent No. 5001368) there is described a configurable logic circuit array having a matrix of discrete sites or
15 cells at each of which is a logic circuit which is adapted to perform a simple logic function. Typically the simple logic function is implemented by means of a two input NAND gate as a primary function and each site has a secondary function. The cells are arranged in
20 tiles being a grouping of 2 x 2 core cells and a selection of different secondary functions are available within a tile. Each cell connects to others by way of a restricted local interconnect structure where each output of a cell connects to a sub-set of the inputs of

its 8 perpendicular neighbours. The cells also connect to each other by medium connection paths that extend between a porting arrangement and at least some of the cells in the zone. The zone connects to the rest of the logic array by means of a porting arrangement which supplies connections between the core cells within the zone via medium buses and with other zones by either global interconnect or by the medium interconnect of the adjacent zones.

10 A restriction of configurable logic circuit arrays is that the full use of all the functions available within the array is not possible, due to restrictions in routing the required signals to and from the function cells. The problem is exacerbated by the requirement to reduce the silicon area used by the device to a minimum, to provide a large function to silicon cost ratio.

20 The logic circuit of the cells provide function resources or can be used as routing resources at a local level. Other interconnect structures must be effectively used and the silicon area used is determined by the functions implemented on the silicon and by the number of programming bits required to control the specified functions. Thus, to keep the silicon cost as low as possible the number of programming bits used needs to be a minimum, with as few

unused programming bits as possible. To facilitate silicon layout of the device, as much uniformity of the physical layout as possible is maintained. In the device according to British Patent Specification No. 9410980.8 all the core cell tiles have the same physical layout and the use of the two input core cell function dictates that implementation of a three input function requires at least two core cells to be used together with the interconnect resources to connect the two cells together. A common function used within circuits is a multiplexer. A two input multiplexer requires three inputs and hence with the prior art multiple core cells have to be used to generate the function, together with the associated interconnect resources used. Making the cells more complex would require a considerable increase in the silicon area.

It is an aim of one aspect of the present invention to increase the number of circuit instances that can be placed in a zone with minimal increase in silicon area. Another aim of the invention is to allow the number of cells that can be used within the zone to be increased with minimal increase in silicon area.

Accordingly, a first aspect of the invention proposes a semi-conductor integrated circuit which, as made, comprises an area thereof formed with a plurality of cells, each cell having a plurality of resources

which comprise storage/programming resources and functional resources, the functional resources comprising at least one of routing resources and logic function resources according to the cell type and
5 whereas, as made, each cell has a finite number of resources in it, and wherein the improvement comprises augmenting the resources available to a cell by utilising resources from at least one other cell.

Thus the present invention proposes sharing of
10 resources between cells so as to maximise silicon utilisation and circuit efficiency.

It is preferred that the cells are disposed in a matrix array which is divided into zones.

In one embodiment, for a cell comprising a two
15 input device which has a primary function and a secondary function, the secondary function can become a three input function (eg. a two input multiplexer) by sourcing the third input from another cell. This is done in such a way that the function of the source cell,
20 which is preferably an adjacent cell, is not disabled. This offers a saving in core cells used, allowing circuits of greater complexity to be implemented upon the array. Its also has the advantage of increasing the speed a function can be run at, when compared with the
25 prior device. More specifically according to this embodiment cells are arranged in tiles, for example

comprising a matrix array of 2×2 core cells, each core cell having a common primary function (eg. 2 input NAND gate) and a secondary function. Usually there will be a number of different secondary functions within a tile, but some core cells may have the same secondary functions. Usually each tile will be identical. The cells are conveniently disposed in rows and columns. For such a construction it is preferred if the third input for the multiplexer is taken from the vertically adjacent cell in the same tile. Also the output is taken from the medium bus output driver. This means that the driving cell can still be configured as a totally independent function, the only restriction being that the driving cell could only output via its local interconnect structure, except where the driving cell also generates the signal to the third input of the driven cell.

The logic sites or cells as made have a local interconnect structure by which any one cell is only connectable with some of the other cells. Usually any one cell will only connect with its eight nearest perpendicular neighbours. In addition, there are medium connection paths by which signals are routed to or from a cell via a port cell associated with the zone in which the cell concerned is located.

In another embodiment, particularly for the

afore-described cells which are constructed identically as regards the number of programming bits for controlling interconnection of cells and where the path for signals between adjacent cells at the zone boundary can be via a medium bus to the port cell and thence to the medium bus of the adjacent zone, there is no need for a local connection path between those adjacent cells. This results in the core cells at the zone boundaries having spare programming bits. For the above-described structure the core cells in the corners of the zones have four spare bits since they could have had four local interconnections that bridged the port cells, ie. two in the horizontal direction, two in the vertical direction. The core cells along the sides have two spare bits corresponding to removal of two local interconnections across the zone boundary. The above-described local interconnects for these boundary core cells can be physically removed.

Greater interconnectivity from the port cells to the core cells would enable designs to be implemented with less routing. This, in turn, would lead to more cells being available for functions rather than routing, thereby increasing the circuit density within the zone. To increase the connectivity at the port cells would require extra programming bits to implement the function. This, in turn, would cause the layout of the

port cell on the silicon to increase in size, leading to an increase in overall device silicon size, and a reduction in function to cost ratio.

According to this aspect of the present invention
5 the programming bits which were previously employed controlling local interconnect at the zone boundaries are reassigned to be used by the port cells to provide greater connectivity from the port cell multiplexers onto the medium busses, thereby enabling connectivity
10 between the global interconnect and the medium interconnect to be increased. The physical location at the reassigned programming bit remains unchanged, ie. they stay in their respective core cells.

Our co-pending published British Patent
15 Application Publication No. 2279168 describes a logic circuit array which, as made, comprises an area thereof formed with a plurality of logic circuits at discrete cells respectively defining a matrix array of cells. The matrix array of cells is sub-divided at least into
20 zones (each comprising a matrix array of cells) and further comprising a porting arrangement for each zone and a hierarchical routing structure comprising global connection paths having selectable connections with a porting arrangement of each zone, medium connection
25 paths extending from the porting arrangement and selectably connectable with at least some of the cells

in a zone, and local direct connection paths comprising a restricted signal translation system.

A restriction of configurable logic circuit arrays is that the full use of all the functions available within the array is not possible, due to restrictions in routing the required signals to and from the function cells across the whole array. The problem is exacerbated when larger devices are implemented, due to the greater number of signals that are required to be passed around the array. As devices become larger, the routing paths across the array becomes slower, due to the physical length of these paths and the loading effects of the resources connecting to them, hence effecting the overall effectiveness of the array.

One solution, and perhaps the obvious solution, is to create more resources at each level, ie. more global connection paths and more medium connection paths. However, this is disadvantageous in terms of silicone utilisation and accordingly the present invention aims to provide a solution.

Accordingly, this aspect of the invention provides a semi-conductor integrated circuit which, as made, comprises an area thereof formed with a plurality of logic circuits and defining a matrix array of cells divided into zones (preferably there is a matrix array of zones) and a zone porting arrangement for each zone,

and further comprising a hierarchical routing structure disposed in successive layers and wherein the successive layers of routing resource comprise: zone interconnect comprising local direct connection paths connecting at least some of the cells with some other cells, and medium connection paths extending from the porting arrangement and selectably connectable with at least some of the cells in a zone; and global connection paths having selectable connections with the porting arrangement of each zone, and further comprising one or more further successive levels of routing resource having a respective porting arrangement to connect that level of routing resource with the preceding level of routing resource.

By adding extra levels of routing hierarchy in this manner, devices larger than that contemplated can be effectively implemented. This allows the architecture to be scaled to any size by the addition of extra levels of routing hierarchy, each separated by a respective porting arrangement. This allows general signals (ie. not clock or reset signals) that are required to be passed across the whole array to use the higher levels of hierarchy before using the lower levels to reach the actual function cells. Whereas the general signals that only require to travel smaller distances are able to do so on the lower levels of the routing

hierarchy. This reduces the loading on each bus since the number of connection points for each bus is reduced. This reduction in loading causes an increase in the speed of the routing paths, hence the speed and effectiveness of the overall array is improved.

For convenience the first further level of routing resource will be referred to as the super-global level and the associated port cells for that level will be referred to as the super-global port cells.

It will be apparent that each porting arrangement connects two adjacent levels of routing hierarchy. Advantageously the porting arrangement also provides connection to other busses within the same layer of hierarchy. The porting arrangement at each level reflects the typical layout of an integrated circuit in which the cells are disposed in rows and columns, and comprises port cells and conveniently there are port cells which extend in the horizontal direction and in the vertical direction. Conveniently these are referred to as horizontal port cells and vertical port cells respectively. Within each layer, there are connection paths, conveniently referred to as busses, which run in the horizontal direction and in the vertical direction. Turns between the horizontal and vertical busses are accommodated by additional busses conveniently referred to as x busses. In the preferred

arrangement these additional busses connect to switches situated at discrete cites within the array, their purpose being to connect the horizontal and vertical segments of these busses together to implement 90° turns within that layer of interconnect. These x busses connect to their level of hierarchy at the porting arrangements. A preferred arrangement is for the switches for these types of busses to be situated within some of the core cells.

Just as there is a port cell for each zone of cells, it is preferred that the porting arrangements for each successive layer will be zoned such that any port cell for that level will have access to a plurality of port cells of the preceding lower level via the connection paths for that level. Thus for example, a super-port cell will connect with a plurality of port cells of the adjacent lower level via the global connection path linking those port cells. Thus, each super-port cell is associated with a plurality of zones. This is conveniently referred to as a super-zone. Usually there will be a decreasing number of port cells in each successively higher layer. Each porting arrangement also allows connection of the same type of busses between adjacent zones in the same layer. Preferably the device has at least three levels of hierarchy. A preferred construction comprises tiles

consisting of four individual core cells, with 25 tiles (100 core cells) making up a zone, and wherein the zones are structured into super-zones with preferably 5 x 5 zones making up a super-zone. There is preferably a matrix of super-zones making up a device with preferably 3 x 3 super-zones plus input/output to make up a device.

One preferred construction of interconnect hierarchy comprises zone interconnect made using either local or medium interconnect, with the port cells (ie. vertical and horizontal port cells) around each zone used to connect the medium busses to the next higher layer of routing hierarchy. Local interconnect is used to connect a cells output for example to each cells nearest perpendicular neighbours, and with the medium busses running horizontally and vertically, providing connections to core cells within the zone. Preferably there are six medium busses per row and column. Interconnect further comprises a global interconnect layer consisting of global and x busses. Connections from zone level medium busses to/from global and x busses are made via the port cells. The global busses, as well as connecting to the port cells, also connect to IO cells and what are termed super-port cells, ie. the port cells of the next higher routing resource. Global busses run vertically and horizontally across the super-zones, ie. between the ports of the next higher level of

resource. The x busses are used to make turns at the global routing level. There are preferably three global busses and one x bus for every row and column of core cells. Switches between horizontal and vertical x busses (to make 90° turn between horizontal and vertical global busses) are situated in each of the core cells. Interconnect further comprises a super-global interconnect layer comprises super-global and super-x busses. This layer is used for connections that span the whole device (that is, across super-zones between the super port cells and other super port cells IO cells). Connections between the global and the super-global resources are made using the super-port cells situated between super-zones. There is preferably one super-global bus and two super-x busses per row and column of core cells. The super-x busses are used to make turns at the super global level. The super-port cells are also used to make liner connections between global and super-global busses.

20 The hierarchical interconnect structure can be further extended to at least a further successive layer of routing conveniently referred to as hyper-global above the super-global interconnect. The super-global interconnect would only span across super-zones whilst the extra layer would extend across at least some super-zones and partitioned by a porting arrangement

(hyper-port cells). The porting arrangement would connect the super-global interconnect to the extra level of interconnect and to the super-global interconnect of the adjoining partitioned area. Within the extra layer, the porting would extend in both horizontal and vertical directions. The busses in the extra layer would run in the horizontal and vertical directions, turns between the horizontal and vertical busses are accommodated by additional busses of type x. These additional busses connect to switches situation at discrete cites within the array, their purpose being to connect the horizontal and vertical segments of these busses together to implement 90° turns within that layer of interconnect, these x busses connect to their level of hierarchy at the porting arrangements.

The various aspects of the present invention will now be described further hereinafter, by way of example only, with reference to the accompanying drawings; in which:-

Figure 1a is a diagrammatic illustration of part of a hierarchical routing structure according to the present invention,

Figure 1b is a diagrammatic illustration of part of a hierarchical routing structure according to another embodiment,

Figure 2.1.A is an overall schematic of an array

structure embodying the invention,

Figure 2.2.A illustrates the location of super-x bus switches,

Figure 2.2.B is a schematic of a super-port cell,

5 Figure 2.3.A illustrates connections for vertical port cells,

Figure 2.3.B is a vertical port cell schematic,

Figure 2.3.C illustrates connections for a horizontal port cell,

10 Figure 2.3.D is a horizontal port cell schematic,

Figure 2.3.E-H are port cell schematics for the top, bottom, left and right edges respectively,

Figure 2.4.A. illustrates diagrammatically a zone structure,

15 Figure 2.4.B illustrates a preferred embodiment of local cell to cell signal connection possibilities,

Figure 2.4.C illustrates the medium bus connections for a Type A tile,

20 Figure 2.4.D illustrates the medium bus connections for a Type B tile,

Figure 2.4.E illustrates a preferred arrangement of tiles in a zone,

Figure 2.5.A is a schematic of a core cell structure,

25 Figure 2.5.B is a diagrammatic illustration of core cell tiling,

Figure 2.5.C is a schematic for the core cell input multiplexer,

Figure 2.5.D is a schematic for the core cell output multiplexer,

5 Figure 2.5.E is a schematic illustration of the combinatorial core cells of Type 1 to 3,

Figure 2.5.F is a schematic for a tile showing the derivation of the third input from vertically adjacent cells,

10 Figure 2.5.G is a schematic for a Type 4 core cell, and

Figure 2.5.I is a diagrammatic illustration showing the arrangement of configuring ram bits for cells in the vicinity of the horizontal and vertical port cell boundary.

15 Figure 1a is a diagrammatic illustration of the present invention. Core cells (CC) comprising a matrix array of integrated circuits are shown in dotted outline as a block of cells. There are three blocks in the illustration. Each block represents a zone (Z). Each zone has associated with it a porting arrangement or port cell (PC). There are discrete local signal connection paths (not illustrated) between inputs and outputs of the core cells, and there are medium signal connection paths (M) which are connectable with the core cells in a zone and extend between the port cells

associated with each zone and the port cell of the next adjacent zone. Only one medium connection path is illustrated, but there will usually be several such paths. In addition there are global signal connection paths (G) which extend across a plurality of zones and are connectable with the port cells of several zones, referred to as super-zones (SZ). There may be a number of global busses.

A further layer of signal interconnect is provided by super-global bus (SGB) which extends across the super-zone (SZ) between super-port cells (SPC) by which signals as the SG bus or busses are routed either between super-zones or down to the global bus(es) (G).

Figure 1b is a diagrammatic illustration of a further embodiment of the invention and extends the arrangement described in Figure 1a to suit yet larger arrays. Like reference numerals denote the same features as described with reference to Figure 1a and the corresponding description applies and has not been repeated. However, for the larger arrays envisage the super-zones SZ are themselves grouped into what are conveniently termed hyper-zones (HyZ), and there are signal connections paths conveniently referred to as hyper-global busses (HyG), which extends across the hyper-zones to form a further layer of routing resource. The hyper-global busses connect with hyper-global port

cells (HyPC) which control passage of the signals between the hyper-zones or down to the lower level resource - namely the super-global busses (SG).

It is preferred that each successively higher level of global routing resource will span a greater number of zones. It is also preferred that each further successive layer of routing resource comprises a decreased number of routing lines, ie. there are fewer super-global busses than global busses, and fewer hyper-global busses than super-global busses, etc.

A preferred array structure is illustrated with reference to Figure 2.1.A for the case of a 22.5k cell array. The array is partitioned hierarchically into super-zones SZ, zones Z and tiles T. Each tile consists of four individual core cells CC, and 25 tiles (or 100 core cells make up a zone. Zones are structured into super-zones, with 9 super-zones plus input/output cells IO making up the 22.5 k cell device. Each zone comprises a matrix array of 10 x 10 core cells and each tile consists of a matrix array of 2 x 2 core cells. A super-zone is made up of a matrix array of 5 x 5 zones. Each zone is bordered on two of its edges by port cells PC conveniently referred to as horizontal port cells HPC along a horizontal edge of the zone and vertical port cells VPC disposed along a vertical edge of the zone. Super-port cells SPC are associated with each of the

super-zones and there are horizontal super-port cells HSPC and vertical super-port cells VSPC. There are 60 input/output cells per side and the corner of the array is occupied by peripheral bus corner switches PBCS. In the illustrated embodiment the array is comprised of a matrix array of 3 x 3 super-zones and separated from each other by respective super-port cells, horizontal and vertical cells as the case may be.

Various levels of signal routing resources are provided and for the preferred embodiment three layers of signal routing resource are employed. The first layer is conveniently referred to as zone interconnect and includes connections within zones using either local or medium interconnect. Local interconnect (see Figure 2.4.B) is used for fast connections to each cell's eight nearest perpendicular neighbours (B, BB, F, FF, U, UU, L, LL), and medium busses (M1-M6) run horizontally and vertically to provide connections to core cells anywhere in the zone. There are six medium busses (M1-M6) per row and column of core cells. Port cells (HPC and VPC) around each zone are used to connect the medium busses to the next layer of routing hierarchy - the global layer described further hereinafter. These port cells also allow medium bus connections to neighboring zones (see Figures 2.3.A-2.3.D).

The second layer of routing, referred to as

global interconnect consists of global G and x busses referenced x. Connections from zone level medium busses to/from global and x busses are made via the port cells PC. Global busses run vertically and horizontally across the super-zones (see Figure 2.4.A). The x busses are used to make turns at the global routing level. There are three global busses and one x bus for every row and column of core cells. Switches between horizontal and vertical x busses (to make 90° turns between horizontal and vertical global busses) are situated in each of the core cells.

Global busses, as well as connecting to port cells, also connect to IO cells and super-port cells.

The third routing layer is conveniently referred to as super-global interconnect and it refers to routes that run across the device between respective super-port cells or between super-zones and/or IO cells) and comprise super-global SG and super-x (ref. SX) busses with each spanning a super-zone. Connections between the global and super-global resources are made using the super-port cells SPC, situated between super-zones (see Figures 2.2.B). There is one super-global bus SG and two super-x (SX) busses per row and column of core cells. The super-port cells are also used to make linear connections between global and super-global busses (ie. from global to global, or super-global to

super-global busses).

In addition, there are special routing resources comprising clocks and resets which are distributed via dedicated interconnect to specified cells in the array.

5 As far as functional resources are concerned, in the illustrated embodiment, all the core cells can be used as a simple NAND gate (see Figure 2.5.A). In addition, each core cell also implements a range of secondary functions. The cells are conveniently referred to as
10 types 1, 2, 3 and 4 (see Figure 2.5.B). Each has the NAND as its primary function, whilst types 1, 2 and 3 have X-OR/MUX as the secondary function and type 4 has as its secondary function a D-type flip-flop-abbreviated DFF (with reset and enable). The clocks and
15 resets are distributed via the dedicated interconnect to the DFF cells in the array. In addition, user configurable IO cells provide connections to/from the array (via super-global, super-x, global, x and medium busses). IO cells are also arranged into zones - the IO
20 zones pitch match the array zones, and there are four IO cells for every 10 core cells. Each IO cell can also connect to a four bit peripheral bus that runs around the array. The peripheral bus can be divided into four sections using switches at the corners of the device.
25 It can be used either as a conventional bus, or as a wired-OR bus for address de-coding etc.

The super-global resources SG, SX provide the top level interconnection. Super-global busses SG provide connections between super-zones, and to/from the IO cells. Super-x busses SX are used to implement 90° turns between super-global busses and are also used to connect to the IO cells. Super-port cells SPC provide connections between the super-global and the global resource layers. There is one super-global bus for every row/column of core cells. Each super-global bus spans a single super-zone, and can be connected only at the super-port cells or in the case of super-zones that are at the edge of the array to/from the IO cells. Super-global busses are a fast method of carrying signal long distances.

There are two super-x busses for every row/column of core cells. Like super-global busses, each super-x bus spans a single super-zone, and can be connected to only at the super-port cells, or (for super-zones at the each of the array) at the IO cells. In addition, super X-busses connect to switches SXBS situated in the core cells so that horizontal and vertical super-x busses can be joined to implement 90° turns. Turn switches are located in every tile (ie. every 4 core cells). Figure 2.2.A shows the detail of the positioning of the switches within each tile. Super-x busses may also be used to route signals across super-zones (ie. like

super-global busses), but they are a slower resource for this than using super-global busses, because of the extra loading of the switches.

Two points should be noted about the super-x bus switches. Firstly, the pattern of super-x bus switches is continuous across zone boundaries. This is simply a feature of the fact that zones are 10 x 10 core cells but the super-x bus pattern repeat every 4 core cells. Secondly, Figure 2.2.A shows the super-x bus switched in core cells 2 and 3. This is logically correct. However, in order to evenly distribute configuration RAM bits between the four core cells in a tile, the switches will physically be placed in core cells 1 and 2.

The device illustration in Figure 2.1.A is partitioned into 9 super-zones. Between each super-zone are horizontal and vertical super-port cells - one super-port cell for every row and column of core cells.

Figure 2.2.B shows the available connections in schematic form for a super-port cell. Horizontal and vertical super-port cells have identical functionality. Each super-port cell requires 24 configuration RAM bits.

Each I/O cell can be connected to/from the super-global bus and both super-x busses in three rows/columns of core cells. Since there are four I/O cells to every zone (ie. 10 core cells), 1 in 5 rows/columns can be connected to either one or two I/O

cells (all other rows/columns to connect to just one I/O cell).

Each array zone (10 x 10 core cells) has ten port cells on each side (see Figure 2.1.A). The port cells above and below a zone are call the Vertical Port Cells (since they handle vertical interconnect, and the port cells to the left and right of each zone are called the Horizontal Port Cells (since they handle the horizontal interconnect). Port cells are used for connections between zonal routing resources (medium busses) and the global resources (global and 'x' busses). Port cells also support connections between medium busses in adjacent zones. All 6 medium busses in each row/column connect to the port cells (M1..6). Primary clocks and resets are distributed via the vertical port cells. Port cells are symmetrical, and buffering is only used when changing from one layer of hierarchy to another.

Figure 2.3.A represents the Vertical Port cell connections, and shows how the zone and global busses are connected. Connections between zones (ie. medium bus to medium bus, or 'x' bus to medium bus) are unbuffered. Port cells are arranged in pairs, aligned with the core cells tiling. This is to allow port cells to access medium busses in both columns of the tile. For example, medium bus M2 of the right hand column in Figure 2.3.A can be connected to medium, global or 'x'

busses of the left hand column using the crossover connection at the bottom of the port cell pair (ie. by using the centre "A2" mux connection in the left hand port cell in Figure 2.3.A). The medium busses from the adjacent columns are called "M1P" or "M2P" in the schematic diagrams. Connections to and from the global routing layers (ie. to and from the global and 'x' busses, via the "A3" and "C3" muxes) are buffered in both direction. The "B" mux in the centre of each port cell allows two direct, unbuffered connections from M3-M3 and M4-M4 in adjacent zones. Four primary clock/reset lines run through the vertical port cells (2 clocks and 2 resets). Alternate port cells have either a primary clock or reset multiplexer - these select the source for the dedicated clock and reset lines to the D flip-flop core cells. A programmable inversion and tie-off are also provided in the clock/reset multiplexer. The primary clock/reset can be connected into the zone, from the clk/rst mux, via muxes "C3" and "C2". Secondary clocks/resets are sourced, via "C3" and the clk/rst mux, from the global busses. Thirty configuration RAM bits are required to programme each vertical port cell. A complete schematic is shown in Figure 2.3.B.

Horizontal Port Cells provide similar resources to the vertical port cells, with two differences.

Firstly, the "B" mux provides a crossover switch between M3 and M4. The "B" mux in the centre of each horizontal port cell allows two direct, unbuffered connections either from M3-M3 and M4-M4, or M3-M4 and
5 M4-M3, in adjacent zones, and secondly there is no clock-reset distribution in the horizontal port cells. Figure 2.3.c shows horizontal port cell connections, and Figure 2.3.D is the full schematic.

Twenty eight configuration RAM bits are required
10 to programme each of the horizontal port cells.

Port cells are symmetrical and are designed to interface to/from zones on either side of the port cell (top and bottom for vertical ports, or left and right for horizontal ports).

15 However, around the edges of the array, port cells only have any array on one side (IO cells are on the other side). Therefore, cut-down port cells are used around the edges of the array, to save configuration RAM bits and to reduce bus loading.

20 There are four different edge port cells, for each side of the array. Schematics for each edge port cells are shown in Figures 2.3.E to 2.3.H.

The bottom, left and right edge port cells need 13 configuration RAM bits for programming. The top-edge
25 port cells requires 17 configuration RAM bits (the extra bits in this cell are needed for clock/reset mux).

Connections should always be made through the port cell multiplexers (for example, global to medium bus through Mux A3 and Mux A2). Connections should not be made in and out of the same side of a port cell mux ("bank-shots"). For example, G1 should not be joined to G2 via Mux A3. The exception to this rule is that medium busses may be joined using Muxes A2 and C2. Connections to 'x' busses are to be buffered. For example, M to X via A1-A2 not allowed, but the preferred route is M to X via A1-A3.

The device in Figure 2.4.A shows a zone structure for the device. The medium, global and 'x' busses are only shown in one row/column in Figure 2.4.A for clarity. The super-global and super-x busses are not shown in Figure 2.4.A because they do not connect to the zone level resources. Local interconnect provides a connection from the output of each core cell to one input multiplexer of its eight nearest perpendicular neighbours. Local interconnect is the fastest connection between core cells. The available connections are shown in Figure 2.4.B. Local interconnect does not continue across zone and super-zone boundaries. The naming convention used in Figure 2.4.B is to name destination core cells relative to the source core cell. For example, core cell UU is two core cells up from the driving cell, B is one core cell back,

and FF is two core cells forward.

There are 6 medium busses per row/column of core cells (M1..6). Medium interconnect is used for connections within a zone that are not possible using local interconnect, and for inter-zone connections via the horizontal and vertical port cells. Four of them can also be connected to global resources via the port cells. M1..2 and M5..6 can all be connected to global or 'X' busses via the port cells. M3..4 do not connect up to global-layer routing, but are used for inter-zone connections (using the 'B' mux in the port cells). See Section 2.3.3 (Port Cells). Inside each zone, the horizontal and vertical medium busses connect to the input and output multiplexers of each core cell. Each core cell has two input multiplexers - each of which can connect to three medium busses. The core cell output multiplexer can drive 6 medium busses. There are 12 medium busses running vertically and horizontally across each core cell. Therefore, each core cell can only connect to half the total number of medium busses available. To compensate for this, there is a hierarchical, repeated pattern of medium bus connections across each zone. M1..4 connections are arranged in a pattern that repeats every tile (2 x 2 core cells) and M5..6 are transposed in adjacent tiles. Each core cell in a tile has the same M5..6 connections,

but adjacent tiles have M5/M6 connections reversed (ie. M5 inputs become M6 inputs, M6 outputs become M5 outputs, and vice versa). Figure 2.4.C shows the medium bus connections in a tile, where M5 is an input and M6 is an output. This is called Tile Type A. Figure 2.4.D shows the medium bus connections for the adjacent tile (M6 inputs and M5 outputs) - called Tile Type B. Figure 2.4.E shows the arrangement of Tiles A and B within a zone. Note that the M1..4 connections are identical in both tile types.

Each IO cell can be connected to/from two of the medium busses (M3/4) in three rows/columns of core cells.

The core cell is the basic functional block in the array. Each core cell can be programmed to implement either a 2-input NAND gate, or one of alternative functions. Each core cell has the same basic structure - shown in Figure 2.5.A.

Each core cell contains:

two input multiplexers (7 inputs each (3, medium, 4 local) plus a tie-low function)
input crossover switch, to swap the selected local interconnect between input muxes
programmable inversion on both inputs
primary function NAND gate
alternative function block (provides 1 or more

additional functions, eg. Mux or DFF)

medium bus output multiplexer (connects to 6
medium busses)

X-bus switch

5 super-X bus switch (core cell types 1 and 2 only)

Core cells are arranged in 2 x 2 tiles - each
tile supports all the available alternative functions.
Three cells in each tile have combinatorial functions
(ie. NAND + XOR and MUX) and the fourth cell in each
10 tile supports a D-type flip-flop (ie. NAND + DFF). The
reset and clock for the DFF cell (type 4) are provided
via dedicated interconnect from the vertical port cells.
Each DFF can have its reset tied inactive if required-
the reset enable is located in the Type 3 cell to
15 provide an even distribution of RAM bits between the 4
core cells. Core cells also contain the SX and X bus
switches. Figure 2.5.B shows the arrangement of cells
within each tile.

Each core cell has two input multiplexers (A, B)
20 - each mux can input from 3 medium busses and 4 local
connections. The local interconnect can be swapped
between the A and B muxes using a crossover switch. See
Figure 2.5.C.

There is a pull-down function on both the A and
25 B input muxes. When either or both of the A and B input
muxes are unused (ie. no inputs selected), their

respective outputs ("A OUT" and/or "B/OUT") are pulled low. The pull down devices are not shown in Figures 2.5.C for clarity.

5 The output multiplexer connects each cell's output to up to 6 medium busses (3 vertical and 3 horizontal). If a core cell is not required to drive a medium bus, then the output mux of that cell may be used to connect medium busses together. The core cell output mux may also be used to supply the 'select' input
10 for the 2-to-1 MUX function in vertically adjacent core cells.

It will be apparent from the above that the core cell provides the basic functional block in the array and that each core cell can be programmed to implement
15 either a 2-input NAND gate, or one of a range of alternative functions. Each core cell has the same basic structure as shown and described with reference to Figure 2.5.A. Reference is also directed to Figures 2.5.B, 2.5.C and 2.5.D. Core cell Types 1 to 3 all have
20 identical functional resources. They differ only in the location of the super x bus switches and reset enable logic. Figure 2.5.E shows the schematic for the Type 1 to 3 core cells - but exclusively the input multiplexer and medium bus output multiplexer. Programmable
25 inversion is provided at both the A and B inputs by 2 RAM bits. Two further ram bits determine the core cell

function. The illustration shows a third input to the core cell at 100 and is the select line for the 2-to-1 mux. It is sourced from the vertically adjacent core cell in the tile (ie. the core cell above for Types 3 and 4, and the core cell for Types 1 and 2) - see Figure 2.5.F. Figure 1.5.D shows how the select line is generated from the medium bus output multiplexer. This arrangement means that there are two ways of routing a select line to the 2-to-2 mux function.

10 The first is direct from horizontal/vertical medium buses. The medium bus output mux 102 of the vertically adjacent core cell in the tile is used as an input multiplexer. The mux select can be sourced from one of the six medium buses that are connected to the output mux. The core cell is still available as a
15 function or routing resource, with the limitation that its output mux is no longer available - ie. the output must go onto local interconnect only.

 Secondly, from local/medium interconnect via a
20 routing cell. If the vertically adjacent core cell is used as a routing cell, then the mux select line can also be sourced from any of the signals available at the core cell A and B input multiplexers. The core cell is no longer able to function as a function or routing
25 resource.

 Without the additional direct interconnection to

the mux select from the output mux of a vertically adjacent cell creation of a three input function could not be achieved in a manner which would allow the cell from which the third input is derived to be used for a
5 totally different function or as a routing function.

Thus, the secondary core cell function of at least some of the cells in a tile can have functions with greater than two inputs, and more specifically, three inputs. More particularly, the three input
10 function is 'a multiplexer' and the third input is the select line for the multiplexer. The third input is driven by a resource from an adjacent cell. The source of the third input is the medium bus output multiplexer, and conveniently from an adjacent cell. This resource
15 can be isolated from the output of the adjacent core cells primary function. The adjacent core cells primary function is thus driven out via local interconnect.

Referring now to Figure 2.5.G, here there is illustrated the schematic for a Type 4 cell which
20 provides the D-type flip-flop function. The DFF is configurable with an asynchronous reset if required, and the clock is sourced from the primary or secondary clocks via the vertical port cells. The schematic excludes the schematic for the input multiplexer and
25 medium bus output multiplexer. Programmable inversion is provided on both the A and B inputs. The D-input to

the DFF has a mux cell in front of it. If the cell is configured as a simple DFF then the 2-to-1 mux can be used as a switch to select whether the D-input is sourced from either the A or B core cell input multiplexers (ie. using the programmable inversion on the select input, in conjunction with the pull down).

As mentioned above, for implementation reasons, port cells have to pitch-match to the array of core cells. That is to say, the vertical port cells have to have the same number of programming bits as the core cells. For the described architecture, each core cell requires 26 configuration RAM bits - these are physically arranged in 2 rows of 13, with the cells logic in between. This core cell layout restricts the dimensions of the port cells. Vertical port cells must use multiples of 13 bits to match to core cells. Twenty-six RAM bits in 2 rows of 13 is a practical maximum when taking die size into account. Horizontal port cells have to pitch match to the number of rows of RAM in a core cell (ie. 2 rows), and there is no theoretical limit on the number of bits in a row. However, the width (ie. bits/row) of the horizontal port cell has to be limited to minimise the total width of the die. For the proposed architecture, 2 rows of 12 bits gives a reasonable width. This is illustrated in Figure 2.3.I.

Eight of the RAM bits of the core cells are normally allocated to enable local interconnections with the cells eight nearest perpendicular neighbours. However, since local interconnect is not continuous across zone and super zone boundaries in this embodiment, core cells around the edge of zones have either 2 or 4 redundant RAM bits which are normally used for local interconnect inputs. Core cells in the corner of a zone have 4 redundant bits, since both vertical and horizontal local interconnect is redundant/removed. Other edge cells only lose 2 bits. These redundant bits are borrowed by the port cells to increase the RAM bits available to each port cell. This is illustrated diagrammatically on Figure 2.3.I. It will be understood that the position of the rearranged redundant bits in Figure 2.3.I is for illustration only as these may not be the actual positions of the bit locations borrowed in the final implementation. However, it will be understood that the borrowed bits physically remain in the core cell and their output is routed to the port cell. The utilisation of the RAM bits in this way allows the port cells to have greater functionality and/or connectivity, thereby increasing the effectiveness of the array. Thus programming bits that are physically placed in one cell but are unused by that cell can be used in an adjacent

cell. The cell from which the programming bit or bits are borrowed can still be configured to function totally independently.

CLAIMS



1. A semi-conductor integrated circuit which, as made, comprises an area thereof formed with a plurality of cells, each cell having a plurality of resources
5 which comprise storage/programming resources and functional resources, the function resources comprising at least one of routing resources and logic function resources according to the cell type, as made each cell has a finite number of resources in it, and wherein the
10 improvement comprises augmenting the resources available to a cell by utilising resources from at least one other cell without preventing the donor cell from functioning independently.
2. A semi-conductor integrated circuit as claimed
15 in claim 1 in which the augmented resources is a routing resource.
3. A semi-conductor integrated circuit as claimed in claim 2 in which the routing resource provides a further input to the cell.
- 20 4. A semi-conductor integrated circuit as claimed in claim 1, 2 or 3 in which the cell whose resources are augmented is a cell with logic circuit capability.
- 25 5. A semi-conductor integrated circuit as claimed in claim 4 in which there are a plurality of logic circuit cells which are arranged in tiles and wherein the cells



of each tile have at least one simple logic function in common and at least one secondary function, and in which there are a number of different secondary functions and each tile has at least one of each different secondary function.

6. A semi-conductor integrated circuit as claimed in claim 5 in which the common logic function is a two input NAND function.

7. A semi-conductor integrated circuit as claimed in claim 5 or 6 in which the secondary function is a function with greater than two inputs.

8. A semi-conductor integrated circuit as claimed in claim 7 in which the secondary function has three inputs.

9. A semi-conductor integrated circuit as claimed in claim 7 or 8 in which the secondary function is a two to one multiplexer.

10. A semi-conductor integrated circuit as claimed in any one of claims 2 to 9 in which the cell providing the augmented resource can still carry out totally independent functions.

11. A semi-conductor integrated circuit as claimed in any one of claims 5 to 10 in which the routing resource comprises a signal path extending between two cells in the same tile and column from a medium bus output multiplexer of one cell to one of three input gates of



two a input multiplexer of the other cell.

12. A semi-conductor integrated circuit as claimed in claim 1 in which the augmented resource is a storage/programming resource.

5 13. A semi-conductor integrated circuit as claimed in claim 12 in which a programming bit unused by one cell is used by an adjacent cell.

14. A semi-conductor integrated circuit as claimed in claim 12 or 13 in which the majority of cells comprise
10 logic cells disposed in a matrix array and divided into a plurality of zones with each zone being bounded by port cells which provided routing functions to connect the logic cells with medium or global signal path using their programming bits, the logic cells having
15 programming bits to selectably control connection with a restricted number of other cells via direct (local) connection paths and wherein the port cell receives the augmented resource from a logic cell.

15. A semi-conductor integrated circuit as claimed in
20 claim 14 in which the port cell receives two programming bits from the adjacent core logic cell.

16. A semi-conductor integrated circuit as claimed in claim 14 in which there are horizontal and vertical port cells and a logic cell in a corner of the zone which is
25 adjacent to both the horizontal and vertical port cells donates four programming bits, two to each of the



adjacent horizontal and vertical port cells.

17. A semi-conductor integrated circuit which, as made, comprises an area thereof formed with a plurality of logic circuits and defining a matrix array of cells
5 divided into zones, there being a plurality of zones in the array and a zone porting arrangement for each zone, and further comprising an hierarchical routing structure disposed in successive layers and wherein the successive layers of routing resource comprises:-

10 zone interconnect comprising local direct connection paths connecting at least some of the cells with some other cells, and medium connection paths extending from the zone porting arrangement and selectably connectable with at least some of the cells
15 in a zone; and

global connection paths having selectable connections with the zone porting arrangement of each zone; and further comprising

20 one or more further successive levels of routing resource having a respective porting arrangement to connect that level of routing resource with the preceding level of routing resource.

18. A semi-conductor integrated circuit as claimed in claim 17 in which each successive level of routing
25 resource spans a greater number of zones.

19. A semi-conductor integrated circuit as claimed



in claim 17 or 18 in which each further successive level of global routing resource comprises a decreased number of routing lines.

20. A semi-conductor integrated circuit as claimed in
5 claim 17, 18 or 19 in which there are a matrix array of zones.

21. A semi-conductor integrated circuit as claimed in
any one of claims 17 to 20 in which for an arrangement
where the cells are disposed in rows and columns, the
10 zone porting arrangement comprises horizontal and
vertical port cells which extend in the horizontal and
vertical directions respectively.

22. A semi-conductor integrated circuit as claimed in
claim 21 in which each level of routing resource has one
15 or more porting arrangements and each porting
arrangement has access to a plurality of the port cells
of the preceding level via the signal connection paths
for that level of routing resource.

23. A semi-conductor integrated circuit as claimed
20 in any one of claims 17 to 22 in which there are a
plurality of signal connection paths to each of the
further successive levels of routing resource.

24. A semi-conductor integrated circuit as claimed in
claim 23 when appendant to claim 22 in which there are a
25 plurality of signal connection paths which extend in the
horizontal direction and a plurality of signal



connection paths which extend in the vertical direction.

25. A semi-conductor integrated circuit as claimed in claim 24 in which there are a plurality of horizontal and vertical signal connection paths for each row and
5 column of cells.

26. A semi-conductor integrated circuit as claimed in claim 25 in which for each cell there is at least one horizontal and vertical signal connection path which is connectable.

10 27. A semi-conductor integrated circuit as claimed in any one of claims 17 to 26 in which there are a plurality of matrix arrays of zones, each defining a super-zone, each zone having a plurality of global connection paths communicating with the zone porting
15 arrangements for each zone and extending across the super-zone between the porting arrangements for a first level of routing resource, hereinafter referred to as the super-global level, said first level of routing resource being selectably connectable with the porting
20 arrangement for a plurality of super-zones.

28. A semi-conductor integrated circuit as claimed in claim 27 in which the cells are grouped in tiles with each tile comprising a two by two matrix array of cells and each zone comprises a matrix array of tiles.

25 29. A semi-conductor integrated circuit as claimed in claim 27 or 28 in which the super-port cell for each



super-zone provides selectable connections between the super-global connections spanning adjacent super-zones or from super-global connections to the global connections for that super-zone.

- 5 30. A semi-conductor integrated circuit constructed and arranged substantially as hereinbefore described with reference to and as illustrated in the accompanying drawings of Figures 1a or 1b or 2.5.G or 2.5.I, or any
10 one or more of these or the other drawings in combination.